

Sub-GeV dark matter as pseudo-Goldstone from the seesaw scale

Michele Frigerio^{a,b}, Thomas Hambye^c, Eduard Masso^{d*}

^a *CNRS, Laboratoire Charles Coulomb, UMR 5221, F-34095 Montpellier, FRANCE*

Université Montpellier 2, Laboratoire Charles Coulomb, UMR 5221, F-34095 Montpellier, FRANCE

^b *Institut de Física d'Altes Energies, Universitat Autònoma de Barcelona, E-08193 Bellaterra, SPAIN*

^c *Service de Physique Théorique, Université Libre de Bruxelles, 1050 Bruxelles, BELGIUM*

^d *Grup de Física Teòrica, Departament de Física, Universitat Autònoma de Barcelona, E-08193 Bellaterra, SPAIN*

Abstract

Pseudo Nambu-Goldstone bosons (pNGBs) are naturally light spin-zero particles, which can be interesting dark matter (DM) candidates. We study the phenomenology of a pNGB θ associated with an approximate symmetry of the neutrino seesaw sector. A small coupling of θ to the Higgs boson is induced radiatively by the neutrino Yukawa couplings. By virtue of this Higgs portal interaction (i) the pNGB acquires a mass m_θ proportional to the electroweak scale, and (ii) the observed DM relic density can be generated by the freeze-in of θ -particles with mass $m_\theta \simeq 3$ MeV. Alternatively, the coupling of θ to heavy sterile neutrinos can account for the DM relic density, in the window $1 \text{ keV} \lesssim m_\theta \lesssim 3 \text{ MeV}$. The decays of θ into light fermions are suppressed by the seesaw scale, making such pNGB sufficiently stable to play the role of DM.

*E-mails: frigerio@ifae.es, thambye@ulb.ac.be

1 Introduction

The astrophysical and cosmological evidence for dark matter (DM) requires that a new particle is introduced, with constrained mass and interactions. These interactions must in particular maintain the DM candidate stable on cosmological time scales and generate the observed DM relic abundance. Many candidates can satisfy all constraints. Much fewer, though, provide a convincing explanation for the required values of the DM parameters. As examples, in most models (i) an ad-hoc symmetry is postulated to insure stability, and/or (ii) the DM mass is assumed to be of order (or below) the electroweak (EW) scale as demanded by the Weakly Interacting Massive Particle paradigm, and/or (iii) new interactions with strength similar to the EW ones, not motivated by other means, are just postulated. A scenario which is potentially predictive, and could provide motivations for some of these assumptions, consists in identifying the DM candidate with the pseudo-Nambu-Goldstone boson (pNGB) of a spontaneously broken global symmetry.

An interesting feature of such scenario is that the DM couplings to the Standard Model (SM) fields, which could render the DM candidate unstable on cosmological times, are suppressed by powers of the spontaneous symmetry breaking (SSB) scale. As a result, provided this scale is large, the decay width of the particle can be sufficiently small. A well-known example is the axion, the pNGB associated with the Peccei-Quinn $U(1)$ symmetry, broken spontaneously at the scale f_{PQ} . The axion is a viable DM candidate when f_{PQ} lies in the interval $\sim (10^{12} - 10^{14})$ GeV. Another possibility is to consider SSB at the seesaw scale, which should be large to account for the smallness of the neutrino masses. One global symmetry related to the seesaw is $U(1)_{B-L}$. In the case that it is spontaneously broken by an EW singlet, coupled to heavy sterile (right-handed) neutrinos, the corresponding NGB is known as the singlet Majoron [1]. Models with a pseudo-Majoron as DM have been studied [2]-[7]. More generally, we will consider any (family-dependent) global symmetry whose spontaneous breaking contributes to sterile neutrino masses. Symmetries associated with the seesaw sector have the advantage that the pNGB interactions are anticipated to exist anyway to account for neutrino masses.

Another remarkable property of the pNGB scenario, relevant for the DM phenomenology, is that the mass and the scalar potential interactions of a pNGB vanish in the limit of exact symmetry. Therefore, when the symmetry is explicitly broken by a unique (or dominant) term, the scalar interactions and the DM mass are necessarily related, and, if the relic density is essentially determined by such interactions, this implies a one-to-one correspondence between the DM relic density and mass. For instance we will consider a framework where the source of explicit breaking generates radiatively a ‘Higgs portal’ interaction of the form $\lambda\theta^2 H^\dagger H/2$, with θ the pNGB particle and H the SM Higgs doublet. This coupling generates a mass $m_\theta^2 = \lambda v^2$ with $v = 174$ GeV and, at the same time, a relic density from the $\theta - H$ interactions. We will show how, given the known thermal distribution of SM particles in the early universe thermal bath, the Higgs portal can generate a DM relic density through the freeze-in mechanism (see e.g. Ref. [8]). The observed relic density is obtained for a unique possible value of the DM mass, which turns out to be $m_\theta \simeq 2.8$ MeV for a Higgs boson mass of 120 GeV. There is also a second possibility to generate the correct relic density, by the freeze-in of the interaction of θ with the sterile neutrinos. These production mechanisms

differ from those previously considered for a pNGB candidate for DM.

A third interesting feature of the pNGB setup is that the DM mass is related to the scale of explicit symmetry breaking, which can be identified with a physical mass scale already present in the theory, analogously to the QCD scale for the axion. For instance, in the scenario we will consider below, the neutrino Dirac Yukawa couplings provide the source of explicit breaking of the global symmetry. Therefore, the DM mass is proportional to the EW scale, $m_\theta^2 = \lambda v^2$, and this provides a justification for the presence of a scalar DM particle at or below the EW scale. Moreover, the coupling λ can be computed as a function of the seesaw couplings. In turn, we will show that in this framework the DM mass can be protected from quadratically divergent corrections. The mechanism we invoke to remove these corrections is a variation of the one proposed long ago by Hill and Ross for very light pNGBs with both scalar and pseudo-scalar couplings [9]. We will consider a $U(1)_X$ symmetry such that its explicit breaking requires several Yukawa couplings at the same time, thus lowering the degree of divergence of the contributions to the pNGB effective potential. This mechanism has similarities with the collective breaking mechanism introduced to protect the EW scale in little Higgs models, where the SM Higgs is the pNGB (see Ref. [10] for a review).

All in all, we will show that the seesaw interactions (i) are associated with global symmetries broken spontaneously at the heavy neutrino mass scale, whose largeness guarantees the stability of the pNGB DM candidate, and (ii) have a source of explicit symmetry breaking built-in, and therefore induce the mass and scalar interactions of the DM. Moreover, they can lead to (iii) a one-to-one correspondence between the DM mass and the DM relic density, with (iv) a justification of the presence of DM at low scale, and with (v) a DM mass protected from large radiative corrections.

The paper is organized as follows. In section 2 we display the relevant effective interactions of our pNGB candidate for DM. In section 3 we show how a pNGB relic density of the order of the observed DM density can be obtained through freeze-in. In section 4 we present a class of seesaw models where the required pNGB interactions are generated naturally. In section 5 we compute the pNGB couplings to SM particles and analyze the corresponding constraints on the DM stability. We conclude in section 6.

2 Effective interactions of the pNGB

In this section we introduce the effective lagrangian of a pNGB θ , associated with a global symmetry of the neutrino sector, broken spontaneously at the seesaw scale f . We postpone to section 4 the detailed description of a class of models where such a light pNGB may emerge naturally. For the present purposes, it suffices to specify the effective interactions of θ at low energy (below f).

In the class of models under consideration, θ couples to a sterile neutrino ν^c as follows:

$$\mathcal{L}_N = -\frac{gf}{2\sqrt{2}}\nu^c\nu^ce^{i\theta/f} + \text{h.c.} = -\frac{m_N}{2}\bar{N}N + \frac{ig}{2\sqrt{2}}\theta\bar{N}\gamma_5 N + \frac{g}{4\sqrt{2}f}\theta^2\bar{N}N + \dots, \quad (1)$$

where g is a Yukawa coupling and the 4-component Majorana neutrino is defined as usual by $N \equiv (\nu^c \ \nu^{c\dagger})^T$. For simplicity, we consider only one sterile neutrino, whose mass $m_N =$

$gf/\sqrt{2}$ is generated by SSB. In this case the lagrangian shown in Eq. (1) is the same as for the singlet Majoron model [1], where the spontaneously broken global symmetry is lepton number. However, we will show in section 4 that in realistic cases with two (or more) sterile neutrinos the relevant symmetry is not the ordinary lepton number, rather each sterile neutrino carries a different charge, and consequently the pNGB θ cannot be identified with the Majoron. We remark that the $\theta - N$ interaction is just an example of a NGB interaction with a heavy fermion charged under the associated global symmetry. As a consequence, the implications for dark matter phenomenology, discussed in the following sections, hold qualitatively in a more general set of theories, with heavy fermions other than the sterile neutrinos.

A non-vanishing potential for the NGB is generated when one introduces a source of explicit breaking of the global symmetry. In the models we will consider, this source is given by a certain set of neutrino Dirac Yukawa couplings, so that the θ effective potential necessarily involves the SM Higgs doublet H . The relevant, radiatively induced coupling is given by

$$\mathcal{L}_H = -\frac{\lambda}{2}\theta^2 H^\dagger H = -\frac{m_\theta^2}{2}\theta^2 - \frac{\lambda v}{\sqrt{2}}\theta^2 h - \frac{\lambda}{4}\theta^2 h^2, \quad (2)$$

where $v \simeq 174$ GeV is the Higgs vacuum expectation value (vev), h is the physical Higgs boson and we adopted the unitary gauge. Note that the mass of the pNGB θ is proportional to the EW vev, $m_\theta^2 = \lambda v^2$. This relation is sometimes referred to as the ‘conformal’ limit, since θ has no bare mass term, rather the Higgs portal interaction λ generates both m_θ and the $\theta - h$ couplings. We remark that the phenomenology (e.g. the DM mass, its relic density, etc.) induced by the Higgs portal in Eq. (2) does not depend on the details of the underlying physics which generates radiatively λ .

In the specific seesaw models that we will build, λ is generated by logarithmically divergent neutrino loops, involving both the $\theta - N$ coupling g as well as neutrino Dirac Yukawa couplings, denoted generically with y . It turns out that the relation between the mass of θ and the neutrino parameters can be written schematically as

$$m_\theta^2 = \lambda v^2 \simeq g^2 y^2 v^2 \frac{\log(\Lambda^2/m_N^2)}{8\pi^2}, \quad (3)$$

where Λ is some cutoff scale at or above f . The light neutrino mass scale is given by the standard seesaw relation, $m_\nu = y^2 v^2 / m_N$. It is useful to use Eq. (3) to express the coupling g as a function of the relevant energy scales in the theory:

$$g = 10^{-3} \left(\frac{m_\theta}{\text{MeV}} \right) \left(\frac{\text{eV}}{m_\nu} \right)^{1/2} \left(\frac{10^9 \text{GeV}}{m_N} \right)^{1/2} \left(\frac{8\pi^2}{\log(\Lambda^2/m_N^2)} \right)^{1/2}. \quad (4)$$

We will take the factor $k \equiv \log(\Lambda^2/m_N^2)/(8\pi^2)$ to be of order one. Then, for a fixed value of m_ν , m_N and also $f = \sqrt{2}m_N/g$ are determined as a function of m_θ and g , which are basically the only two free parameters. We require the scale f to be smaller than the Planck scale, $M_P = 1.22 \times 10^{19}$ GeV. In addition, we assume for simplicity that f is larger than one TeV, in order for the effective theory to contain only the SM particles, the pNGB and the sterile neutrino, while the degrees of freedom involved in the SSB are decoupled.

These two requirements exclude the green shaded regions in the $(m_\theta - g)$ plane shown in Fig. 1. The dotted lines in Fig. 1 correspond to constant values of the sterile neutrino mass, $m_N = m_\theta^2/(g^2 m_\nu k) = 10^2, 10^6, 10^{10}, 10^{14}$ GeV, from top to bottom.

3 Relic density from freeze-in of the pNGB

We now study the role of the interactions in Eqs. (1) and (2) for the production of a θ relic density, which can play the role of DM. In this section, we implicitly assume that θ is stable on cosmological time scales. The region of parameters where this stability is achieved will be analyzed in detail in section 5. There are several mechanisms to produce the DM relic density. We begin by briefly recalling these various possibilities and confront them with the parameter space of our scenario.

A well-known way to produce a cosmological density of a DM particle species is the usual freeze-out: it requires an annihilation cross section large enough to overcome the expansion of the universe and thermalize the DM particle. The interactions of θ come from the terms in Eq. (1), which we refer to as the ‘sterile neutrino portal’, and the terms in Eq. (2), that is, the ‘Higgs portal’. These interactions may or may not lead to the thermalization of θ in the early universe, depending on whether the interaction rate Γ gets larger than the Hubble rate H . If it does, later on θ necessarily decouples from the plasma and, if stable, acts as DM with a certain energy density ρ_θ , which depends in general on its mass and annihilation cross section.

Before determining ρ_θ numerically, it is instructive to identify the thermalization region in first approximation, by evaluating the θ interaction rate Γ and requiring $\Gamma > H$. We start with the Higgs portal, considering for simplicity the $h \rightarrow \theta\theta$ decay only, because in our scenario decays turn out to dominate with respect to scattering processes (see Fig. 2 below). The decay rate is

$$\Gamma(h \rightarrow \theta\theta) = \frac{1}{16\pi} \lambda^2 \frac{v^2}{m_h} \sqrt{1 - \frac{4m_\theta^2}{m_h^2}}. \quad (5)$$

This should be compared to the Hubble parameter in the radiation epoch, which is given by $H = 1.66 \sqrt{g_*^p} T^2/M_P$, with $g_*^p = 106.75$ the degrees of freedom corresponding to the SM content (we shall assume such value in the rest of the paper). Taking $T \simeq m_h$ to be the relevant temperature, we find that, for a coupling

$$\lambda \gtrsim 6 \times 10^{-8} \left(\frac{m_h}{120 \text{ GeV}} \right)^{3/2}, \quad (6)$$

the decays and inverse-decays are able to produce a thermal population of θ -particles. In other words, given the $m_\theta^2 = \lambda v^2$ relation, one needs $m_\theta > 44$ MeV, which translates into the thin vertical line in Fig. 1.¹

¹Notice that Eq. (6) is the condition for DM to thermalize with the SM sector at $T \simeq m_h$ (where the rates reach the maximum, see Fig. 2 below). This condition is different from the one usually considered for the freeze-out scenario, where one requires that the scattering processes are in thermal equilibrium at $T \lesssim m_\theta$. The latter condition leads to a value of λ about two orders of magnitude larger than the one in Eq. (6).

Similarly for the interaction with the sterile neutrino N , the process $NN \rightarrow \theta\theta$ has a rate

$$\Gamma = \langle \sigma v_{rel} \rangle n_N \simeq \frac{g^4}{256\pi} \frac{n_N}{m_N^2} \simeq 2 \times 10^{-4} g^4 m_N, \quad (7)$$

with n_N the sterile neutrino number density, and we used $T \simeq m_N$. This leads to thermal production of θ -particles provided

$$g \gtrsim 10^{-2} \left(\frac{\text{eV}}{m_\nu k} \right)^{1/6} \left(\frac{m_\theta}{\text{MeV}} \right)^{1/3}, \quad (8)$$

where we made use of Eq. (4). This gives the thin diagonal line in Fig. 1. We do not discuss the effects of the couplings of θ to light fermions, because in the models under consideration they turn out to be far too small to thermalize θ (see section 5).

Outside the thermalization region, where Γ is smaller than H , one may still have a final Ω_θ matching Ω_{DM} (as usual we define $\Omega_i \equiv \rho_i/\rho_{crit}$, i.e. the ratio between the energy density ρ_i and the critical density of the universe). The idea is that in the early universe a population of θ -particles is produced through pair production scattering processes, but with a density smaller than the thermal one. For appropriate values of m_θ one may still reach a relic density equal to Ω_{DM} . This process, called freeze-in, actually happens in our scenario as we shall describe below in detail (for generic properties of the freeze-in mechanism see Ref. [8]).

In the rest of this section we first introduce the ingredients to make a numerical calculation of the θ relic density, and we then apply them to the freeze-out and freeze-in processes. Afterwards we also discuss some non-thermal production mechanisms and end up with a summary of our results.

3.1 Boltzmann equation and reaction densities

From the Higgs portal, θ -particles are created from the following scattering processes:²

$$\sigma(hh \rightarrow \theta\theta) \simeq \frac{\lambda^2}{32\pi s} \left(\frac{s - 4m_\theta^2}{s - 4m_h^2} \right)^{1/2} \left(\frac{s + 2m_h^2}{s - m_h^2} \right)^2, \quad (9)$$

$$\sigma(WW \rightarrow \theta\theta) = \frac{1}{9} \frac{\lambda^2}{32\pi s} \left(\frac{s - 4m_\theta^2}{s - 4m_W^2} \right)^{1/2} \frac{s^2 - 4sm_W^2 + 12m_W^4}{(s - m_h^2)^2 + m_h^2 \Gamma_h^2}, \quad (10)$$

$$\sigma(ZZ \rightarrow \theta\theta) = \frac{1}{9} \frac{\lambda^2}{32\pi s} \left(\frac{s - 4m_\theta^2}{s - 4m_Z^2} \right)^{1/2} \frac{s^2 - 4sm_Z^2 + 12m_Z^4}{(s - m_h^2)^2 + m_h^2 \Gamma_h^2}, \quad (11)$$

$$\sigma(f\bar{f} \rightarrow \theta\theta) = \frac{1}{4} \frac{1}{n_c} \frac{\lambda^2}{16\pi s} \frac{m_f^2 (s - 4m_f^2)^{1/2} (s - 4m_\theta^2)^{1/2}}{(s - m_h^2)^2 + m_h^2 \Gamma_h^2}. \quad (12)$$

where Γ_h is the total Higgs decay width and n_c is the number of colours of the fermion f .

²In Eqs. (9)-(13) we included the usual initial-state spin and color average factor, as well as the final state identical particle factor 1/2. In $\sigma(hh \rightarrow \theta\theta)$ we neglected subleading terms of order λ^3 and λ^4 (which are displayed e.g. in Ref. [11]).

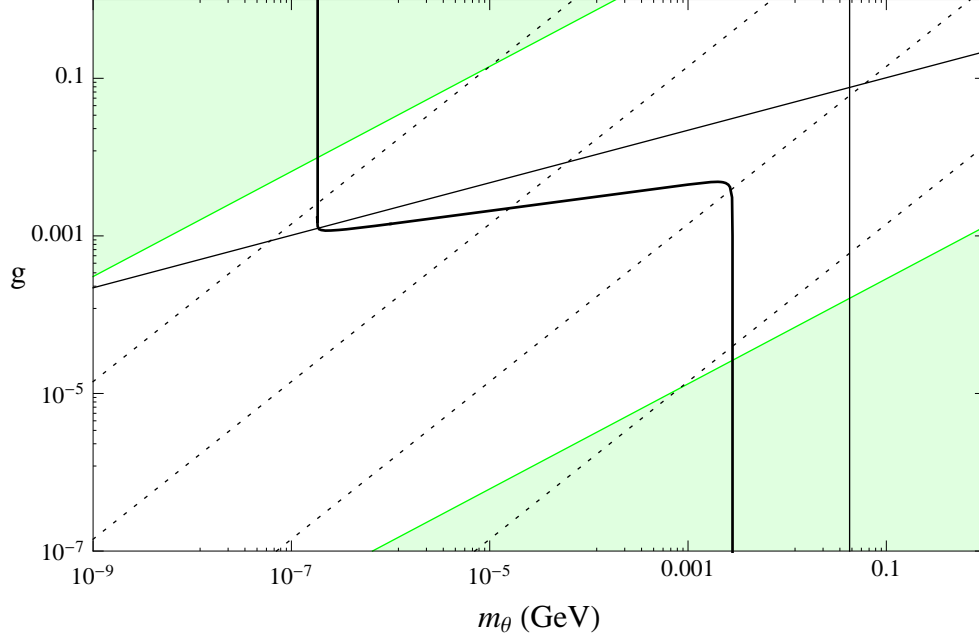


Figure 1: The pNGB coupling to the sterile neutrino, g , versus the pNGB mass, m_θ , in GeV. We fixed $m_\nu = 0.05$ eV and $k = 1$. The upper (lower) green shaded region is excluded by requiring the SSB scale f to be larger than one TeV (smaller than the Planck scale). The four dotted lines correspond to constant values of the sterile neutrino mass: from top to bottom, $m_N = 10^2, 10^6, 10^{10}, 10^{14}$ GeV. The thin diagonal (vertical) line indicates the lower value of g (m_θ) that leads to a thermalization of θ . The thick line corresponds to a θ relic density equal to the observed DM relic density, as follows from the numerical solution of the Boltzmann equation (for $m_h = 120$ GeV). Since θ is produced by freeze-in, its relic density grows with its couplings g and $\lambda = m_\theta^2/v^2$, therefore the region below and to the left (above and to the right) of the thick line corresponds to $\Omega_\theta < \Omega_{DM} (> \Omega_{DM})$.

Similarly, θ -particles can be created through the scattering with a sterile neutrino,

$$\sigma(NN \rightarrow \theta\theta) \simeq \frac{g^4}{256\pi} \frac{1}{\beta_N^2 s} \left(\beta_N \frac{s}{m_N^2} + 2 \log \frac{1 - \beta_N}{1 + \beta_N} \right), \quad (13)$$

where we neglected m_θ in front of m_N and \sqrt{s} and we defined $\beta_N = (1 - 4m_N^2/s)^{1/2}$.

In the following it will be convenient to consider separately the effect of the Higgs decay, Eq. (5), which means that, in order to avoid double counting, the on-shell part of the scattering processes must be subtracted (whenever there is one). In the narrow width approximation this means to perform in Eqs. (10)-(12) the substitution

$$\frac{m_h \Gamma_h}{(s - m_h^2)^2 + m_h^2 \Gamma_h^2} \rightarrow \frac{m_h \Gamma_h}{(s - m_h^2)^2 + m_h^2 \Gamma_h^2} - \pi \delta(s - m_h^2) \theta(s - 4m_i^2), \quad (14)$$

for $i = W, Z, f$. In the following, all scattering quantities we will consider are meant to be the subtracted ones, using Eq. (14). To calculate the relic density that one obtains from

these processes, either through freeze-in or freeze-out, one has to integrate the Boltzmann equation,

$$zH(z)s(z)Y'_\theta(z) = \left[1 - \left(\frac{Y_\theta(z)}{Y_\theta^{\text{eq}}(z)}\right)^2\right] (\gamma_{D_h}(z) + \gamma_{\text{annih}}(z)), \quad (15)$$

with $z \equiv m_h/T$ conventionally taken as the evolution parameter. Here $H(z)$ is the Hubble parameter, $s(z)$ the entropy density, and $Y_\theta \equiv n_\theta/s$ with n_θ the number density. The reaction density γ_{D_h} contains the effect of the Higgs boson decay, while γ_{annih} is the sum of the (subtracted) reaction densities of the scattering processes in Eqs. (9)-(13). The reaction densities are given by

$$\begin{aligned} \gamma_{D_h} &= N_\theta \int d\bar{p}_h f_h^{\text{eq}} \iint d\bar{p}_1 d\bar{p}_2 (2\pi)^4 \delta^4(p_h - p_1 - p_2) |\mathcal{M}|^2 \\ &= N_\theta n_h^{\text{eq}} \frac{K_1(z)}{K_2(z)} \Gamma(h \rightarrow \theta\theta), \end{aligned} \quad (16)$$

$$\begin{aligned} \gamma(a b \leftrightarrow 1 2) &= N_\theta \iint d\bar{p}_a d\bar{p}_b f_a^{\text{eq}} f_b^{\text{eq}} \iint d\bar{p}_1 d\bar{p}_2 (2\pi)^4 \delta^4(p_a + p_b - p_1 - p_2) |\mathcal{M}|^2 \\ &= N_\theta \frac{T}{64 \pi^4} \int_{s_{\min}}^\infty ds \sqrt{s} \hat{\sigma}(s) K_1\left(\frac{\sqrt{s}}{T}\right), \end{aligned} \quad (17)$$

involving the Bessel functions $K_{1,2}$. We have defined $d\bar{p} \equiv d^3p/((2\pi)^3 2E)$. Here $N_\theta = 2$ is the number of θ particles produced per decay or annihilation, $f_i^{\text{eq}} = (e^{E_i/T} \pm 1)^{-1} \simeq e^{-E_i/T}$ is the Maxwell-Boltzmann energy distribution, $|\mathcal{M}|^2$ is the amplitude squared summed over initial and final spins (with no averaging), $s_{\min} = \max[(m_a + m_b)^2, (m_1 + m_2)^2]$, and the reduced cross section is defined by

$$\hat{\sigma}(a b \leftrightarrow 1 2) = \frac{g_a g_b}{c_{ab}} \frac{2[(s - m_a^2 - m_b^2)^2 - 4m_a^2 m_b^2]}{s} \sigma(a b \rightarrow 1 2), \quad (18)$$

with σ the particle physics cross section of Eqs. (9)-(13), $g_{a,b}$ the number of degrees of freedom of the particles a, b and c_{ab} a combinatorial factor equal to 2 (1) if a and b are identical (different).

3.2 Freeze-out

As it is well-known, if θ freezes out relativistically, that is at $T \gtrsim m_\theta$, its relic density is independent of the annihilation cross section, because it decouples when the thermal number density, $Y_\theta^{\text{eq}} \equiv n_\theta^{\text{eq}}/s$ is still independent of the temperature. The formula for the relative density Ω_θ is simply

$$\Omega_\theta h_0^2 \simeq 78 \frac{1}{g_*^s} \frac{m_\theta}{\text{keV}}, \quad (19)$$

with g_*^s the number of degrees of freedom contributing to entropy at θ decoupling. This should match the value

$$\Omega_{DM} h_0^2 = 0.11 \pm 0.01. \quad (20)$$

In these equations, h_0 is the reduced Hubble constant $h_0 = H_0/(100 \text{ Mpc km s}^{-1}) \simeq 0.70$. As a result, the observed relic density can be obtained only for one DM mass value: $m_\theta \simeq 0.15 \text{ keV}$, where we took $g_*^s = 106.75$, valid for a decoupling temperature above 100 GeV. Such a light θ can be only thermalized by the interaction with sterile neutrinos, which decouple at a temperature just below m_N , leading to relativistic θ 's. This result corresponds to the left vertical branch of the thick curve in Fig. 1, which is obtained by integrating numerically the Boltzmann equation.

If instead θ is heavier and thermalizes, to have a small enough relic density requires that it decouples non-relativistically. This can occur through the Higgs portal interaction. In this case the relic density essentially only depends on the annihilation cross section which must take the value $\langle \sigma v_{rel} \rangle \simeq 10^{-26} \text{ cm}^3 \text{s}^{-1}$. To get a sufficiently small relic density, one needs to go to much higher values of m_θ , where the coupling $\lambda = m_\theta^2/v^2$ becomes large enough. The observed relic density is obtained for a unique value of m_θ , because of the relation between m_θ and λ , which is a consequence of the pNGB nature of θ . The freeze-out through the Higgs portal was considered in Ref. [12], assuming this restrictive relation, and it was found that the DM relic density is obtained with $m_\theta = 50 - 70 \text{ GeV}$ (for $m_h = 120 - 180 \text{ GeV}$).

However, as we will see in section 5, the freeze-out values, $m_\theta = 0.15 \text{ keV}$ as well as $m_\theta = 50 - 70 \text{ GeV}$, require values of parameters which are excluded by the instability of θ on cosmological time scales. Therefore in our scenario both relativistic and non-relativistic freeze-out is not a viable way to produce DM. Instead, the freeze-in process turns out to be efficient in an allowed region of the parameter space, as we now discuss.

3.3 Freeze-in

In the freeze-in process θ is produced by the annihilation or decay of a heavier particle X , with rates small enough not to thermalize, leading to a less-than-thermal DM population, which reaches a plateau at $T \sim m_X$ because at smaller T the number of X -particles is Boltzmann suppressed. If the X -particles are the SM particles, as it is the case through the Higgs portal, their number densities are known as a function of the temperature (simply given by their thermal distribution down to a temperature well below their mass). As a consequence, the created relic density depends only on the magnitude of the portal.³

Let us first consider a value of the coupling g sufficiently small to neglect θ production at the seesaw scale (we will see that this is the case for $g \lesssim 10^{-3} - 10^{-4}$, depending on m_θ). In this case only the Higgs portal interactions can account for the observed relic density. As said above, this mechanism is extremely predictive since, given the ‘conformal’ relation $m_\theta^2 = \lambda v^2$, the rates and hence the relic density depend only on the parameter m_θ .

Integrating numerically Eq. (15), it turns out that the observed relic density in Eq. (20) is obtained for

$$m_\theta = (2.76 - 2.86) \text{ MeV} , \quad (21)$$

where we took $m_h = 120 \text{ GeV}$ for definiteness. This result corresponds to the right vertical branch of the thick curve in Fig. 1.

³For production of DM through the Higgs portal along the freeze-in scenario, see also Ref. [13].

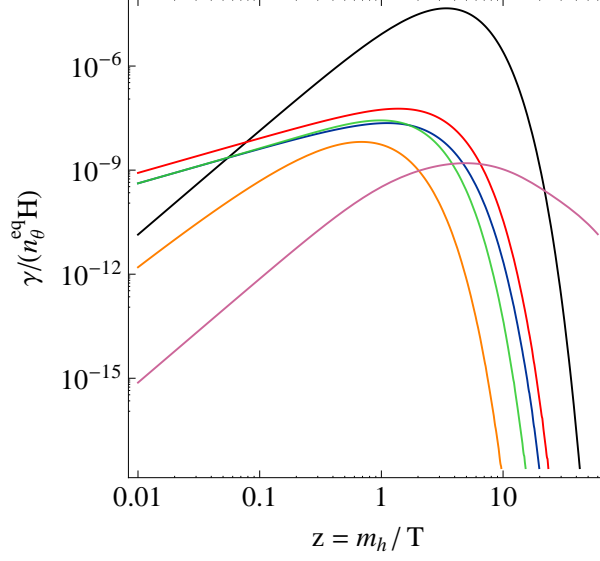


Figure 2: The decay thermalization rate, $\gamma_{Dh}/(n_\theta^{eq}H)$ (black), compared to the various scattering thermalization rates, $\gamma_{annih}^a/(n_\theta^{eq}H)$ for $a = W, Z, h, t, b$ (in red, blue, green, orange and purple respectively), as a function of $z = m_h/T$ and for $m_h = 120$ GeV, $m_\theta = 2.8$ MeV. For freeze-in these rates remain always well below one, which is the thermalization value.

An important difference between freeze-in and freeze-out is that, while the latter is generally dominated by the off-shell annihilation, the former is naturally dominated by the on-shell annihilation, i.e. by the decay of the mediator involved in the annihilation. For a freeze-out, unless the mediator has a mass $\sim 2m_{DM}$, the decay has a small effect because, when the off-shell annihilation goes out-of-equilibrium, it has already decoupled (being more Boltzmann suppressed). For a freeze-in instead, the DM production occurs mostly at T of order the mediator mass, or slightly below, where the Boltzmann suppression is still mild. As a result the decay, which involves less (small) couplings, is naturally dominant, $\gamma_{Dh} > \gamma_{annih}$. This is shown in Fig. 2 where the various reaction densities are plotted for $m_h = 120$ GeV and $m_\theta = 2.8$ MeV. In the approximation where only the decay is included, the Boltzmann equation can be integrated analytically, giving [8]

$$Y_\theta \simeq N_\theta \frac{135 g_h}{8\pi^3 (1.66) g_*^s \sqrt{g_*^p}} \frac{M_P \Gamma(h \rightarrow \theta\theta)}{m_h^2}, \quad (22)$$

where $g_h = 1$. This analytic result turns out to differ from the numerical result by less than 1%.

To understand Eq. (22) let us first discuss why the freeze-in is infrared dominated, with most of the θ production occurring at $T \lesssim m_h$ (just before n_h^{eq} gets too much Boltzmann suppressed). The θ number density produced at temperature T is essentially given by the number of decays per unit time occurring at this temperature, $n_h^{eq}(m_h/T)\Gamma(h \rightarrow \theta\theta)$, times the number of θ -particles produced per decay, N_θ , times the Hubble time, $1/H$. Therefore, taking into account thermal averaging properly, one has $Y_\theta \simeq \gamma_{Dh}/(Hs)$, which is nothing

but the rate of thermalization $\gamma_{Dh}/(n_\theta^{eq}H)$ shown in Fig. 2, times the relativistic value of Y_θ^{eq} . Thus one obtains $Y_\theta \propto M_P m_h \Gamma(h \rightarrow \theta\theta)/T^3$, which increases as T decreases.

Hence one finds $Y_\theta = c[N_\theta n_h^{eq} \Gamma(h \rightarrow \theta\theta)/(Hs)]_{T=m_h}$, with c a numerical factor which turns out to be equal to $(3\pi/2)/K_2(1) \simeq 2.9$. It is larger than one because, as Fig. 2 shows, the maximum value of the decay reaction density occurs at $T \sim m_h/3.5$ rather than at $T \sim m_h$.⁴ With this in mind, and taking into account the dependence $\Gamma(h \rightarrow \theta\theta) \propto \lambda^2/m_h \propto m_\theta^4/m_h$, the DM relic density from the Higgs decay scales as m_θ^5/m_h^3 , so that the value of m_θ needed varies very little with the exact value of the relic density, and the larger the Higgs mass is, the larger m_θ has to be. For example for $m_h = 140, 180, 300$ GeV one gets the observed relic density for $m_\theta = 3.0, 3.6, 4.9$ MeV instead of 2.8 MeV in Eq. (21).

It is interesting to discuss what is the contribution of each scattering channel separately, even though they have a small effect with respect to the decay. As can be seen from Fig. 2, among the scattering channels, the W one gives a contribution slightly larger than the Z and h ones, and much larger than the b and t ones. Numerically, for $m_h = 120$ GeV, the W, Z, h, b, t scatterings contribute to the total number of θ -particles produced in the proportions $1 : 0.4 : 0.4 : 4 \times 10^{-3} : 10^{-1}$ respectively, while the W scattering alone gives a contribution 500 times smaller than the decay channel.

Similarly to the decay, at a temperature $T \gtrsim m_a$, a scattering process $aa \rightarrow \theta\theta$ gives a number density $Y_\theta^a \simeq \gamma_{annih}^a/(Hs)$, which is infrared dominated too. Again, this is nothing but the thermalization rate $\gamma_{annih}^a/(n_\theta^{eq}H)$, displayed in Fig. 2, up to the multiplicative constant Y_θ^{eq} . One gets $Y_\theta^a = c_a[\gamma_{annih}^a/(Hs)]_{T=m_a}$ with c_a a numerical factor which, unlike for the decay, does not take a unique value. It depends on m_θ as well as on the s dependence of the scattering cross section (which fixes the position of the peak of the reaction densities in Fig. 2). For the W, Z and h channels $c_a \simeq 2.2$, whereas for the fermion channels $c_t \simeq 1.7$ and $c_b \simeq 6$. One understands consequently why the W, Z and h channels, which give contributions $Y_\theta \propto M_P/m_{W,Z,h}$, dominate over the b channel, which gives $Y_\theta \propto m_b^3 M_P/m_h^4$. As for the top channel, it is suppressed too because for $s \gtrsim 4m_t^2$ it is more threshold suppressed than the other channels, see Eqs. (9)-(12), and for higher s it gets quickly suppressed by the m_t^2/s^2 asymptotic behaviour of the top cross section (as compared to the $1/s$ asymptotic behaviour of the W, Z and h cross sections).

Let us now move to the region $m_\theta < 3$ MeV, where the Higgs portal alone would lead to a too small relic density. In this case the N annihilation process can do the job, as long as the reheating temperature is large enough to produce a thermal population of sterile neutrinos, which we assume.⁵ The annihilation rate depends on two free parameters only, g and m_θ , with the sterile neutrino mass m_N determined by Eq. (4). From the freeze-in of

⁴Actually if one takes the value of $\gamma_{Dh}/(Hs)$ at its maximum value, i.e. at $T \sim m_h/3.5$, one obtains a number density equal to 70% of the exact result. Therefore, the production of a particle through freeze-in of a decay is essentially given by the maximum value of the corresponding thermalization rate. Note also that if one takes instead $N_\theta n_h^{eq} \Gamma(h \rightarrow \theta\theta)/(Hs)$ at its maximum value, which occurs for $T \sim m_h/3.0$, it turns out that one gets an even better approximation, as this overestimates the exact result by only 3%.

⁵It is well-known that, if a sterile neutrino has Yukawa couplings inducing a light neutrino mass of the order of the atmospheric or solar neutrino mass scales, the decays/inverse decays of the sterile neutrino are in thermal equilibrium at $T \sim m_N$. There are various cosmological issues which can put constraints on thermal sterile neutrinos, see e.g. [14]. However they become relevant for sterile neutrino masses below the \sim GeV scale, which are not pertinent in our case, see Figs. 3,4 below.

the reaction $NN \rightarrow \theta\theta$ one obtains as usual $Y_\theta \simeq [\gamma_{annih}/(Hs)]_{T=m_N}$. One then finds that the DM relic density scales as $g^6 m_\nu/m_\theta$, as confirmed by the numerical integration of the Boltzmann equation.

Our numerical result in the $(m_\theta - g)$ plane is shown in Fig. 1 for $m_\nu = 0.05$ eV. The diagonal branch of the black thick line corresponds to a freeze-in through the sterile neutrino portal, with the correct value of the DM relic density. The required value of the $\theta-N$ coupling is given approximatively by $g \simeq 2 \cdot 10^{-3} (m_\theta/\text{MeV})^{1/6} (\text{eV}/m_\nu)^{1/6}$. Above (below) the line the relic density is too large (small). Since the relic density is proportional to the sixth power of g , the predicted value of g depends very weakly on the cutoff Λ , which we fixed in all the numerical analysis, $k = \log(\Lambda^2/m_N^2)/(8\pi^2) = 1$.

3.4 Non-thermal production

In general, one may think of non-thermal, more model dependent, θ production mechanisms. A class of mechanisms is related to the phase transition at the scale of SSB f . At this epoch (part of) the energy stored in the scalar potential false vacuum may be transferred to the pNGBs. There could also be cosmic strings produced at the phase transition, which would decay into pNGBs, but their exact density depends on model details. Here we assume that these contributions to Ω_θ are negligible. This is the case, in particular, if the universe underwent an inflationary phase at temperatures smaller than f (note that the reheating temperature can still be larger than $m_N = gf/\sqrt{2}$, since the coupling g is much smaller than one in our scenario).

A potential source of non-thermal production of pNGB dark matter could be provided by the oscillations of the field θ around the minimum of its effective potential, in case the value of θ at high temperature is displaced from such minimum, a mechanism well-studied in the case of the axion. Let us summarize the axion case. The axion mass is very suppressed above Λ_{QCD} , therefore at high temperatures the axion behaves as an exact NGB. In this case, at the end of the associated Peccei-Quinn phase transition at scale f_{PQ} , the value of the axion field a can lie in any of the equivalent vacua described by $0 \leq a/f_{PQ} < 2\pi$. Later, when $m_a(T)$ becomes important with respect to the Hubble parameter $H(T)$, the field a begins to oscillate around zero, producing a coherent state of particles at rest with an associated Ω_a^{oscill} (for a review see Ref. [15]). One may naively think that this picture applies also to our scenario, with a negligible $m_\theta(T)$ at high temperature (where the EW symmetry is restored), and the oscillations beginning only at $T \sim \text{TeV}$, when m_θ is generated. If this were the case, taking the values of f and m_θ in the range we considered for the freeze-in production, one would conclude that Ω_θ^{oscill} overcloses the universe, unless the initial value of θ is tuned to be much smaller than f , or inflation takes place at temperatures below the TeV scale.

However, the temperature dependence of $m_\theta(T)$ is very different from the one of the axion, whose mass is generated non-perturbatively by an anomaly at Λ_{QCD} . In our case, m_θ is generated, instead, by an explicit breaking of $U(1)_X$ and therefore it does not vanish at high temperatures. On the contrary, we expect $m_\theta(T)$ to receive large thermal corrections. Even if the EW symmetry is restored for $T \gtrsim \text{TeV}$, there are contributions to m_θ that are not proportional to the Higgs vev v , such as $\delta m_\theta^2 \sim g^2 y^2 \Lambda^2 / (16\pi^2)^2$ (see the discussion at the end of section 4). These lead to $m_\theta^2(T) \simeq \text{constant} \cdot T^2$ at high temperature. A quantitative

estimate of a possible non-thermal production of θ would then require a non-trivial study of the thermal evolution of the field θ during and after the phase transition at scale f . Still, we just notice that $m_\theta(T)$ is typically larger than $H(T)$ already at $T \sim f$. Thus, we argue that the field θ does not acquire a random initial value of order f , but rather it sits in the minimum of the potential already at high temperature, or in other words coherent oscillations are strongly suppressed. On the basis of this argument, we assume that Ω_θ^{oscill} is negligible, and that the thermal freeze-in dominates the DM relic density.

3.5 Summary

Our results on the θ relic density are summarized in Fig. 1, where we show the line $\Omega_\theta = \Omega_{DM}$ in the $(m_\theta - g)$ plane. The behaviour of the line with the correct relic density can be easily understood. For $m_\theta \simeq 0.15$ keV, g must be larger than (or equal to) the value needed to thermalize θ . For progressively larger m_θ , the value of g should correspond to less and less thermalization. For $0.15 \text{ keV} < m_\theta < 3 \text{ MeV}$, one can see that the required coupling g is progressively smaller than the thermalization value, indicated by the thin diagonal line. When one approaches $m_\theta \simeq 3 \text{ MeV}$, g becomes rapidly smaller as the Higgs portal begins to contribute to the θ number density. Once the Higgs portal produces by itself the correct amount of θ -particles, g must be small enough to make the sterile neutrino freeze-in negligible.

We conclude that the observed DM relic density can be obtained thermally for

$$0.15 \text{ keV} \lesssim m_\theta \lesssim 3 \text{ MeV} , \quad (23)$$

where the exact upper value depends on the Higgs mass, as discussed in section 3.3. Note that this upper bound holds as an absolute prediction as soon as the Higgs portal dominates the DM production. This is necessarily the case, in particular, if the reheating temperature lies below the sterile neutrino mass scale. More generally, the prediction $m_\theta \simeq 3 \text{ MeV}$ holds for any pNGB whose Higgs portal interaction λ gives the dominant contribution to its mass and relic density, independently of the associated global symmetry and of the source of explicit symmetry breaking which induces λ . In particular the SSB scale may be different from the seesaw scale.

Of course the prediction for m_θ would change if the number of DM particles was not negligible already at temperatures higher than the EW scale. In this case one has to produce less of them through freeze-in, which means that m_θ has to be smaller. In other words, $m_\theta \simeq 3 \text{ MeV}$ constitutes an absolute upper bound for the freeze-in mechanism, and it holds as soon as the freeze-in production dominates over the initial population.

4 Approximate symmetries of the seesaw sector

In this section we consider the SM augmented with sterile neutrinos ν_i^c , with a global symmetry broken spontaneously at the seesaw scale, and we discuss in some detail the generation of a mass for the associated pNGB.

Let us consider the most general Yukawa interactions to be added to the SM in the presence of gauge singlet fermions,

$$-\mathcal{L}_{\nu^c} = l_\alpha m_{\alpha j} \nu_j^c \left(\frac{H}{v} \right) + \frac{1}{2} \nu_i^c M_{ij} \nu_j^c + \text{h.c.} , \quad (24)$$

where l_α are the lepton doublets and H is the Higgs doublet, whose neutral component acquires a vev $v = 174$ GeV. Here and in the following, the mass parameters $m_{\alpha j}$ and M_{ij} are intended in a generalized sense as dynamical scalar fields that may or may not acquire a non-zero vev. Thus, the above lagrangian has a global $U(1)_L$ symmetry with the following lepton number assignments:

$$L(l) = +1 , \quad L(\nu^c) = -1 , \quad L(H) = L(m) = 0 , \quad L(M) = +2 , \quad (25)$$

where we dropped flavour indices. When M acquires a vev, $U(1)_L$ is spontaneously broken and a massless NGB appears in the spectrum of the theory, the singlet Majoron [1]. One can write more explicitly

$$M_{ij} = g_{ij} \Phi , \quad \Phi \equiv \frac{\rho}{\sqrt{2}} e^{i\theta/f} , \quad (26)$$

where Φ is a complex scalar with $L = 2$, the vev $\langle \rho \rangle = f$ breaks spontaneously lepton number, and θ is the Majoron.

There are various possible sources of $U(1)_L$ explicit breaking coming from other sectors of the lagrangian. There may be soft terms in the scalar potential involving Φ which break lepton number; their mass scale, whose size is arbitrary, determines the induced Majoron mass (see e.g. [7]). Another possible source of $U(1)_L$ explicit breaking are quantum gravity effects at the Planck scale M_P , which can break in general all non-gauge symmetries. Assuming these effects are suppressed by powers of M_P , they can be used to generate a Majoron mass at the keV scale [2, 3], which has been extensively studied as DM candidate [3, 4, 5, 6].

In this paper we will consider a different source of explicit breaking of global symmetries, provided by the set of the Yukawa couplings [9]. This allows to relate the size of the pNGB parameters to the fermion mass scales already present in the theory. We will focus on symmetries of the lagrangian in Eq. (24) other than $U(1)_L$, with each lepton carrying in general a different charge. In this case the symmetries are respected only by some matrix elements $m_{\alpha j}$ and M_{ij} , and they are explicitly broken if some other matrix elements are non-zero. The mass and couplings of the pNGBs will be completely determined by the seesaw parameters and by the choice of the cutoff, since they can receive cut-off dependent quantum corrections.

The pNGB mass is, in general, quadratically sensitive to the cut-off. In order to understand the origin of quadratic divergences and the mechanism to remove them, consider first the explicit breaking of a $U(1)$ symmetry in a theory with only one sterile neutrino ν^c :

$$-\mathcal{L}_{\nu^c} = \frac{1}{2} \nu^c (M_a e^{i\theta/f} + M_b) \nu^c + \text{h.c.} = \frac{1}{2} (M_a + M_b) \bar{N} N - \frac{i M_a}{2f} \theta \bar{N} \gamma_5 N - \frac{M_a}{4f^2} \theta^2 \bar{N} N + \mathcal{O}((\theta/f)^3) , \quad (27)$$

where $M_{a,b}$ are real mass parameters. Here M_a is generated by the spontaneous breaking of the $U(1)$ symmetry associated with the NGB θ , while M_b breaks this symmetry explicitly. It

is instructive to compute the fermion loops generating the pNGB mass term, $m_\theta^2 \theta^2/2$, using the right-hand side of Eq. (27): if M_b is zero, the two relevant one-loop diagrams cancel each other, as expected for an exact NGB. However, in the presence of the explicit breaking, a non-zero quadratically divergent contribution is left, which is given by

$$m_\theta^2 \sim \frac{1}{8\pi^2} \frac{M_a M_b}{f^2} \Lambda^2. \quad (28)$$

The effective theory below the scale f may still contain a light scalar θ , since m_θ can be parametrically small, but its NGB nature is obscured by the quadratic dependence on the details of the ultraviolet completion.

In the presence of more than one fermion family one can define several family-dependent $U(1)$ symmetries. In general, they are explicitly broken by some mass matrix entries. However, certain $U(1)$'s will be broken only when several entries are non-zero at the same time: this is the key to reduce the degree of divergence of the radiative contribution to m_θ , where θ 's are the associated pNGBs [9]. This fact can be understood in the language of the effective potential V_{eff} , considering the full (active + sterile) neutrino mass matrix \mathcal{M} : the term in V_{eff} quadratic in the mass matrix, $\sim \text{Tr}(\mathcal{M}\mathcal{M}^\dagger)\Lambda^2$, is invariant under certain $U(1)$ symmetries, that is, it does not depend on the associated pNGBs. Thus, the potential of these pNGBs contains at most terms quartic in the mass matrix, $\sim \text{Tr}(\mathcal{M}\mathcal{M}^\dagger\mathcal{M}\mathcal{M}^\dagger) \log \Lambda^2$, which are only logarithmically sensitive to the cutoff (for an application of this idea to eV scale sterile neutrinos, with the pNGB playing the role of dark energy, see [16]). Of course non-abelian symmetries are also possible, with several NGBs and, potentially, qualitatively different phenomena, but this extension will not be needed for our purposes and will not be considered in this paper.

In order to identify the combination of matrix entries that induces a pNGB mass, we rewrite Eq. (24) with the replacement $l_\alpha(H/v) \rightarrow \nu_\alpha$, in the minimal case of two sterile neutrinos, which are sufficient for realistic light neutrino masses:

$$-\mathcal{L}_{\nu^c} = \nu_\alpha (m_{\alpha 1} \ m_{\alpha 2}) \begin{pmatrix} \nu_1^c \\ \nu_2^c \end{pmatrix} + \frac{1}{2} (\nu_1^c \ \nu_2^c) \begin{pmatrix} M_{11} & M_{12} \\ M_{12} & M_{22} \end{pmatrix} \begin{pmatrix} \nu_1^c \\ \nu_2^c \end{pmatrix} + \text{h.c.} \quad (29)$$

Here a sum over $n_f = 3$ active flavours is understood ($\alpha = e, \mu, \tau$). Let us identify the matrix entries whose phases can be removed. There are $2n_f + 3$ mass terms, which in general have a phase. One can absorb $n_f + 2$ of these phases by redefining the $n_f + 2$ neutrino fields present in Eq. (29). As a result, there remain $n_f + 1$ complex matrix entries in Eq. (29) (either n_f phases in m and one in M , or equivalently $n_f - 1$ phases in m and two in M). When one sets these $n_f + 1$ entries to zero, there are no physical phases left.

Now, suppose these $n_f + 1$ complex entries vanish and the non-zero entries are generated by a set of scalar fields $\Phi_a \equiv \rho_a \times \exp(i\theta_a/f_a)/\sqrt{2}$, acquiring a vev $\langle \rho_a \rangle = f_a$. Then, by an appropriate redefinition of the lepton fields, one can remove all the phase fields θ_a from the Yukawa lagrangian of Eq. (29). Thus, these fields will have only derivative couplings to the leptons, resulting from the redefined lepton kinetic terms. This means that the lagrangian has exact $U(1)$ symmetries, broken spontaneously by $\langle \Phi \rangle$'s, and θ_a are exact NGBs.

When one of the zero matrix entries is switched on, however, it is no longer possible to remove all the phases. This means that one $U(1)$ is explicitly broken, with the associated

pNGB θ acquiring non-derivative couplings. Thus, a non-zero m_θ is generated by neutrino one-loop diagrams. To estimate m_θ , consider the set of the $(n_f + 3)$ non-zero matrix entries. One can check that θ could be rotated away from Eq. (29) if any out of four of these entries were switched off. This means that a $U(1)$ symmetry is recovered when any of them is put to zero, and therefore m_θ must be proportional to the product of the four entries. As a consequence, quadratically divergent contributions to m_θ turn out to be absent.

The θ mass can be controlled by the entries of M only, by those of m only, or by both. Let us specify three models with a $U(1)_X$ symmetry, which are representative of these three generic possibilities:

- (i) The symmetry is broken explicitly in the singlet neutrino sector. Two independent entries of the Majorana mass matrix M are allowed by $U(1)_X$, while the third is forbidden. For example take $X(\nu_1^c) = 0$, $X(\nu_2^c) = 1$ and one scalar field Φ with $X(\Phi) = -2$. Then M_{11} is allowed and M_{22} is generated by the vev of Φ , while $M_{12} = 0$. The NGB coupling to neutrinos reads $M_{22} \exp(i\theta/f) \nu_2^c \nu_2^c / 2$. (The Dirac neutrino sector is not relevant here: the m entries may or may not respect $U(1)_X$, in any case the leading symmetry breaking effects are controlled by M .) When M_{12} is different from zero, neutrino loops give a non-zero contribution to m_θ as long as they contain four appropriate mass insertions:

$$m_\theta^2 \sim \frac{1}{8\pi^2} \frac{M_{11} M_{12} M_{22} M_{12}}{f^2} \log \frac{\Lambda^2}{\mu^2} , \quad (30)$$

where we included a loop suppression factor and μ is the renormalization scale, which can be taken of the order of the sterile neutrino masses. The cancellation of quadratic divergences as well as the structure of the non-vanishing contribution to m_θ are dictated by our symmetry argument, and they can be both explicitly checked by computing the loops.

- (ii) The symmetry is broken explicitly within the Dirac neutrino sector only. The entries on a given row of the Dirac mass matrix m are both allowed, while in the other rows only one entry is allowed. For example take $X(\nu_e) = 1$, $X(\nu_{\mu,\tau}) = 3$, $X(\nu_1^c) = -1$, $X(\nu_2^c) = 1$ and a scalar field ϕ with $X(\phi) = -2$. Then m_{e1} is allowed, m_{e2} , $m_{\mu 1}$ and $m_{\tau 1}$ are generated by the vev of ϕ , and $m_{\mu 2} = m_{\tau 2} = 0$. (For simplicity we suppose that the scalar fields ϕ and Φ that can contribute to m and to M belong to different sets. The singlet neutrino sector is assumed to respect $U(1)_X$: in the present example M_{12} is allowed, thus giving an equal mass to the two sterile neutrinos.) This case reproduces closely the scenario discussed by Hill and Ross [9], which dealt with the quark Dirac mass matrix. Once the entries explicitly breaking $U(1)_X$ are switched on in m , one obtains

$$m_\theta^2 \sim \frac{1}{8\pi^2} \frac{m_{\alpha 1} m_{\alpha 2} m_{\beta 2} m_{\beta 1}}{f^2} \log \frac{\Lambda^2}{\mu^2} , \quad \alpha \neq \beta . \quad (31)$$

- (iii) The symmetry is broken explicitly only by the interplay of the Dirac and Majorana neutrino mass matrices. One possibility is that $U(1)_X$ allows for all the entries of

M , e.g. taking the charges $X(\nu_1^c) = -1$, $X(\nu_2^c) = 1$ and $X(\Phi) = 2$; this symmetry can allow only one between $m_{\alpha 1}$ and $m_{\alpha 2}$, depending on the charge of ν_α . Another possibility is that $U(1)_X$ allows for all the entries of m , but it forbids two entries of M , e.g. taking the charges $X(\nu_\alpha) = X(\nu_1^c) = 0$ and $X(\nu_2^c) = -X(\phi) \neq 0$. As usual the mass of the NGB θ is generated when one zero entry is switched on. One finds

$$m_\theta^2 \sim \frac{1}{8\pi^2} \frac{M_{ij} M_{kl} m_{\alpha 1} m_{\alpha 2}}{f^2} \log \frac{\Lambda^2}{\mu^2}, \quad ij \neq kl. \quad (32)$$

In order to roughly quantify the energy scales under discussion, let us indicate with $M = gf$ ($m = yv$) a generic seesaw (electroweak) scale, possibly suppressed by a small Yukawa coupling g (y) with respect to the $U(1)_X$ SSB scale f (the EW scale v), and let us introduce also the neutrino mass scale $m_\nu = m^2/M$. Up to a loop suppression, in the three cases m_θ is of order (i) $(M/f)M = gM$, at or below the seesaw scale; (ii) $(m/f)m = gm_\nu$, at or below the neutrino mass scale; (iii) $(M/f)m = gm$, at or below the electroweak scale. In the following we will concentrate on the last possibility, since we are interested in scalar DM candidates below the electroweak scale.

Before moving to the other properties of the pNGB relevant for DM, some comments are in order to assess the soundness of the above estimate for the pNGB mass m_θ . We have seen that, in general, in the presence of explicit breaking of the associated global symmetry, m_θ can be sensitive quadratically to the cutoff of the theory. However, if the global symmetry is broken only by the contemporary presence of several couplings, then m_θ shall be proportional to the product of all these couplings. In this case the Feynman diagrams contributing to the pNGB mass will have a lower degree of divergence and thus quadratic divergences vanish. This possibility was often employed in model-building in the past, and recently was extensively used in the context of little Higgs models, under the name of “collective breaking” (see Ref. [10] for reviews). The motivation is to stabilize the electroweak scale against large quantum corrections. In these models the Higgs is a NGB of a symmetry broken spontaneously at the scale $f_{ew} \sim \text{TeV}$, whose mass is generated at one-loop level by two couplings $y_{1,2}$, whose contemporary presence breaks explicitly the symmetry. As a consequence, $m_h^2 \sim y_1^2 y_2^2 f_{ew}^2 \log(\Lambda^2/\mu^2)/(16\pi^2)$, that is, the sensitivity to the cut-off Λ is only logarithmic. However, in general, the quadratic divergence reappears at the two-loop level, with a correction to the Higgs mass $\delta m_h^2 \sim y_1^2 y_2^2 \Lambda^2/(16\pi^2)^2$, which indicates that Λ cannot be larger than $\sim 4\pi f_{ew}$, in order for the theory to remain natural.

It is instructive to compare this little Higgs scenario with our scenario, where the DM candidate θ is a pNGB associated with SSB at the scale f , of the order of the seesaw scale. In this case the pNGB mass generated at one-loop can be written schematically as $m_\theta^2 \sim g^2 y^2 v^2 \log(\Lambda^2/\mu^2)/(16\pi^2)$. The mass m_θ is not proportional to f , but rather to the electroweak scale v . Besides, it can be much smaller than v because of the four powers of Yukawa couplings and the loop suppression. This is true even with a huge cutoff Λ , since the dependence on it is only logarithmic. However, at higher orders the quadratic divergence may reappear, e.g. through a two-loop diagram with a virtual Higgs exchanged across the neutrino loop, leading to $m_\theta^2 \sim g^2 y^2 \Lambda^2/(16\pi^2)^2$. This result is not surprising, since the mechanism we adopted explains why the pNGB mass is related to the EW scale, but it does not address the stability of the EW scale against radiative corrections. In other words, as

already remarked by Hill and Ross [9], the Higgs sensitivity to quadratic corrections also enters in m_θ at higher order. In order for this two-loop correction to be negligible with respect to the one-loop estimate, one needs the cutoff of the Higgs boson loops, Λ_H , to be smaller than $\sim 4\pi v$. (By enlarging the global symmetry, it may be possible to remove also the two-loop quadratic divergence, see e.g. [16], but in general higher orders reintroduce the problem.) If one wants to stabilize a theory which includes a scale much larger than v (e.g. the scale f), one must address the usual hierarchy problem, e.g. postulating supersymmetry broken at $\Lambda_{susy} \sim \text{TeV}$, or a strongly interacting sector that condenses at $\Lambda_c \sim \text{TeV}$ generating dynamically the EW scale. In this paper we assume the stability of the EW scale is realized, and thus we can adopt our one-loop estimates for the pNGB mass and couplings, in order to study phenomenology.

4.1 A seesaw model leading to the Higgs portal

Let us consider in some detail a specific $U(1)_X$ symmetry of the neutrino sector, such that the pNGB θ acquires radiatively a coupling to the SM Higgs. Take the sterile neutrino charges $X(\nu_1^c) = -1$ and $X(\nu_2^c) = 1$, and a scalar field Φ with charge $X(\Phi) = 2$, whose vev breaks the symmetry spontaneously. The interaction lagrangian involving the sterile neutrinos and the pNGB θ reads

$$-\mathcal{L}_{\nu^c-\theta} = l_\alpha(m_{\alpha 1} \ m_{\alpha 2}) \frac{H}{v} \begin{pmatrix} \nu_1^c \\ \nu_2^c \end{pmatrix} + \frac{1}{2}(\nu_1^c \ \nu_2^c) \begin{pmatrix} M_{11}e^{i\theta/f} & M_{12} \\ M_{12} & M_{22}e^{-i\theta/f} \end{pmatrix} \begin{pmatrix} \nu_1^c \\ \nu_2^c \end{pmatrix} + \text{h.c.} \ . \quad (33)$$

Here M_{12} is a mass term allowed by $U(1)_X$, while M_{11} and M_{22} are generated after SSB. The symmetry is broken explicitly by either $m_{\alpha 1}$ or $m_{\alpha 2}$, depending on the $U(1)_X$ charge assigned to the lepton doublet l_α .⁶ Note that θ can be rotated away by rephasing the neutrino fields if $m_{\alpha 1} = 0$ or $m_{\alpha 2} = 0$. In this case the symmetry is restored and θ is a true NGB. The dependence on θ can be removed also when two independent entries of M are vanishing, therefore the symmetry breaking effects must vanish also in this limit. As we discussed in the previous section, this need of four different couplings to break the symmetry is the key to cancel quadratic divergences.⁷

The interactions in Eq. (33) generate an effective potential for the pNGB θ . We assumed that M_{ij} and $m_{\alpha i}$ are real and positive (see discussion below) and we performed the one-loop computation of the effective potential, which can be written as

$$V_{eff} = \frac{\lambda}{2}\theta\theta H^\dagger H + \mathcal{O}(\theta^4) \ . \quad (34)$$

⁶ The choice of X -charges is made to realized a model of type (iii), in accordance with the classification of the previous section. This choice is not unique. An equivalent possibility is a $U(1)_X$ symmetry with charges $X(\nu_1^c) = -1$, $X(\nu_2^c) = 1$, $X(l_\alpha) = 0$, broken spontaneously by a scalar field of charge $X(\phi) = 1$. Then in Eq. (33) one should replace $m_{\alpha 1} \rightarrow m_{\alpha 1}e^{i\theta/f}$ and $m_{\alpha 2} \rightarrow m_{\alpha 2}e^{-i\theta/f}$, while the entries M_{11} and M_{22} are independent from θ and represent the source of explicit symmetry breaking. It is easy to check that the two realizations are equivalent, since the lagrangian is the same up to a rephasing of the neutrino fields.

⁷Note that the field θ can be rotated away from Eq. (33) by phase redefinitions, as long as the n_f entries that break explicitly $U(1)_X$ are set to zero. This is slightly different from the general case discussed below Eq. (29), where $n_f + 1$ entries must vanish, for the trivial reason that here the same phase field appears (with opposite sign) in two independent entries. The only physically relevant fact still holds: the product of four independent entries is needed to generate a mass for θ .

We find that the quadratically divergent contributions cancel explicitly, as expected, and the logarithmically divergent ones give

$$\lambda \simeq \frac{1}{4\pi^2} \frac{M_{12}(M_{11} + M_{22})}{f^2} \frac{\sum_{\alpha} m_{\alpha 1} m_{\alpha 2}}{v^2} \log \frac{\Lambda^2}{\mu^2}, \quad (35)$$

with a renormalization scale $\mu \sim M_{ij}$ and up to tiny corrections of higher order in $m_{\alpha i}/M_{kl}$. This type of models is particularly predictive, because the pNGB mass is generated by the same loops that generate λ , that is, m_{θ} is obtained from Eq. (34) by replacing the Higgs with its vev, $m_{\theta}^2 = \lambda v^2$. By taking a common value $m_N = gf/\sqrt{2}$ ($m = yv$) for each entry of the Majorana (Dirac) neutrino mass matrix, Eq. (35) reduces to Eq. (3), up to a factor $2n_f$ accounting for the sum over flavour indexes. As we saw in section 3, the value of the coupling λ has a crucial role for the (partial) thermalization of θ and thus for the determination of its relic density.

A comment on the CP symmetry is in order. In general, the mass terms M_{ij} and $m_{\alpha i}$ may be complex, that is, they may carry phases that cannot be removed by a redefinition of the fields. These phases correspond to an explicit breaking of CP. In this case θ would have both scalar and pseudo-scalar couplings to the fermions, with characteristic phenomenological signatures [9]. Moreover, the effective potential would contain terms odd in θ , such as $\mu \theta H^\dagger H$ (with $\mu \ll v$), that induce a vev for θ and a small mixing with the Higgs. This may endanger the stability of θ , since it would couple linearly to SM particles. Still, these couplings can be very small and in addition the θ decays may be kinematically forbidden if m_{θ} is sufficiently small. In this paper we do not investigate this more complicated CP-violating possibility, and we rather assume that CP is a good symmetry of the seesaw sector. In this case M_{ij} and $m_{\alpha i}$ in Eq. (33) are all real and θ preserves its pseudo-scalar nature, since Eq. (33) is invariant under a CP transformation with $\theta \rightarrow -\theta$. Usually the DM stability requires an additional (discrete) symmetry that forbids its decay into SM particles. In the present scenario the CP symmetry itself guarantees that θ couples linearly only to the heavy sterile neutrinos, and we will see in section 5 that this makes θ sufficiently long-lived.

5 Constraints on the pNGB lifetime

We first derive the couplings of θ to the SM fermions and gauge bosons, and compare its decay width into light SM particles with the lifetime of the universe. Subsequently, assuming that θ -particles account for the whole DM density, we discuss the more stringent astrophysical and cosmological constraints which exist on the decay of θ to neutrinos, electrons and photons.

5.1 θ couplings to the SM particles

The pNGB lifetime is determined by its couplings to light fermions. These come from the θ - N interaction, through the ν - N mixing induced by the Dirac neutrino masses. In this way θ can decay to light neutrinos (at tree level) and to charged fermions (at one-loop level). In turn, these couplings to SM fermions could induce, through triangle loop diagrams, couplings to SM gauge bosons. These decays should be sufficiently slow to make the DM lifetime longer

than the age of the universe, $\tau_0 \simeq 5 \cdot 10^{17}$ s. This provides interesting constraints on the seesaw parameters, as we now discuss.

Neutrinos. At tree level, θ couples only to light neutrinos as follows:

$$\mathcal{L}_{\theta\nu\nu} = \frac{i}{2} \sum_{\alpha\beta} \frac{(\mu_\nu)_{\alpha\beta}}{f} \bar{\nu}_\alpha \gamma_5 \nu_\beta \theta, \quad (\mu_\nu)_{\alpha\beta} \equiv - \sum_{ij} X_{ij} (m M^{-1})_{\alpha i} M_{ij} (M^{-1} m^T)_{j\beta}. \quad (36)$$

Here X_{ij} is the power of $e^{i\theta/f}$ associated with the sterile neutrino mass matrix entry M_{ij} . In the singlet Majoron model $X_{ij} = 1$ for all i, j , therefore one finds $\mu_\nu = m_\nu \equiv -m M^{-1} m^T$, which is the usual seesaw formula. On the contrary, in our models based on a family-dependent $U(1)_X$ symmetry, θ couples differently to each entry of M , as for example in the model of Eq. (33). This lagrangian leads to a total decay width (into both neutrinos and antineutrinos)

$$\Gamma(\theta \rightarrow \nu\nu) = \frac{1}{16\pi} g_{\theta\nu\nu}^2 m_\theta, \quad g_{\theta\nu\nu}^2 \equiv \frac{\text{Tr}(\mu_\nu^\dagger \mu_\nu)}{f^2}, \quad (37)$$

where we neglected $(m_{\nu i}/m_\theta)^2$ corrections in the phase space, $m_{\nu i}$ being the light neutrino mass eigenvalues. In the Majoron case one obtains $g_{\theta\nu\nu}^2 = \sum_i m_{\nu i}^2/f^2$. This width is smaller than $1/\tau_0$ for $g_{\theta\nu\nu} \lesssim 3 \times 10^{-19} (\text{MeV}/m_\theta)^{1/2}$. Such tiny coupling is natural for θ , because the pNGB couplings are suppressed by the SSB scale f : one needs $f \gtrsim 3 \times 10^9 \text{GeV} (m_\nu/\text{eV}) (m_\theta/\text{MeV})^{1/2}$, where we used the one-family approximation, $g_{\theta\nu\nu} \simeq m_\nu/f$. To translate this bound in the $(m_\theta - g)$ plane we insert the relation $f = \sqrt{2} m_N/g$ in Eq. (4), which gives

$$g_{\theta\nu\nu} \simeq 10^{-21} \left(\frac{\text{MeV}}{m_\theta} \right)^2 \left(\frac{g}{10^{-3}} \right)^3 \left(\frac{m_\nu}{\text{eV}} \right)^2 k. \quad (38)$$

The condition $1/\Gamma(\theta \rightarrow \nu\nu) > \tau_0$ excludes the region above the blue dashed line in Figs. 3,4.

Charged fermions. The pNGB θ couples also to charged fermions, through EW one-loop diagrams. This effect arises because of the mixing between the sterile and the weakly-interacting neutrinos, in particular it is also operative in the singlet Majoron model [1, 17]. The coupling to quarks is induced by a one-loop $\theta - Z$ mixing diagram, with neutrinos in the loop. The coupling to charged leptons is generated by an analog Z -exchange diagram plus a triangle diagram with W -exchange. The resulting coupling can be written in a compact form in the one-family approximation, as follows:

$$\mathcal{L}_{\theta f \bar{f}} = i g_{\theta f \bar{f}} \bar{f} \gamma_5 f \theta, \quad g_{\theta f \bar{f}} = \pm \frac{G_F}{(4\pi)^2} \frac{\sqrt{2} m_N}{f} m_f m_\nu, \quad (39)$$

where the sign is $+$ ($-$) for up quarks and charged leptons (down quarks), and G_F is the Fermi coupling constant. This effective coupling can be suppressed only taking a small $g = \sqrt{2} m_N/f$, since all the other parameters in $g_{\theta f \bar{f}}$ can be determined experimentally. The decay width is given by

$$\Gamma(\theta \rightarrow f \bar{f}) = \frac{1}{8\pi} g_{\theta f \bar{f}}^2 m_\theta \left(1 - \frac{4m_f^2}{m_\theta^2} \right)^{1/2}. \quad (40)$$

We remark that all models where the singlet Majoron is given a mass larger than one MeV and plays the role of DM candidate are constrained (or already excluded) by such decays into charged fermions.

In our scenario, the requirement $1/\Gamma(\theta \rightarrow f\bar{f}) > \tau_0$ leads to an upper bound $g \lesssim 2(\text{MeV}/m_\theta)^{1/2}(\text{eV}/m_\nu)$ (MeV/m_f), where we have assumed $m_\theta \gg 2m_f$ for simplicity. This bound excludes the region above the red dashed line in Fig. 4. As a consequence, taking into account that g and m_θ are related by Eq. (3), and that the Dirac mass $m = yv$ cannot be larger than $\sim \text{TeV}$, we conclude that m_θ close to the EW scale would imply a lifetime shorter than τ_0 . In fact, even stronger constraints on the decay width into charged fermions come from astrophysical and cosmological observations (see section 5.2), which will lead to a stronger constraint, $m_\theta < 1 \text{ GeV}$. In particular, in this class of models θ cannot be the $\sim 50 \text{ GeV}$ DM candidate produced by the freeze-out of the λ interaction with the Higgs, discussed in section 3.2. One is left with the possibility of a sub-GeV DM candidate, because in this case the decays into charged fermion pairs are sufficiently slow (or forbidden kinematically), and the correct relic density can be generated by the freeze-in mechanism.

In the realistic three-family case, the coupling $g_{\theta f\bar{f}}$ is generically of the same order, but with a complicated dependence on flavour parameters. In particular one may argue that some cancellation can take place, to reduce the θ decay width. In addition, we have seen in section 4 that the pNGB mass generation depends crucially on the interplay between the flavour structures of the matrices m and M . Therefore, if one were to study the whole parameter space, one should know the explicit dependence of $g_{\theta f\bar{f}}$ on the mass matrix entries.

For illustration, we display the result for the model in Eq. (33), considering for simplicity only one lepton doublet. Writing the Dirac mass matrix as $m \equiv (m_1 \ m_2)$ and keeping terms up to order $(m_i/M_{jk})^2$, the effective coupling of θ to fermions through the mixing with the Z gauge boson is given by

$$\mathcal{L}_{\theta f\bar{f}}^Z = i \left(\frac{M_{11}}{f} F_1 - \frac{M_{22}}{f} F_2 \right) \frac{2\sqrt{2} G_F}{(4\pi)^2} m_f T_{3f} \bar{f} \gamma_5 f \theta, \quad (41)$$

$$F_1 = m_1^2 \left(c^2 s^2 K + \frac{c^2}{M_1} + \frac{s^2}{M_2} \right) - m_2^2 c^2 s^2 K - m_1 m_2 cs \left((c^2 - s^2) K - \frac{1}{M_1} + \frac{1}{M_2} \right), \quad (42)$$

$$F_2 = m_2^2 \left(c^2 s^2 K + \frac{s^2}{M_1} + \frac{c^2}{M_2} \right) - m_1^2 c^2 s^2 K + m_1 m_2 cs \left((c^2 - s^2) K + \frac{1}{M_1} - \frac{1}{M_2} \right), \quad (43)$$

where T_{3f} is the third isospin component of the left-handed part of the fermion f , $M_{1,2}$ are the eigenvalues of the matrix M , which is diagonalized by a rotation of angle δ , defined by $\tan 2\delta = 2M_{12}/(M_{11} - M_{22})$, and we denoted $c \equiv \cos \delta$ and $s \equiv \sin \delta$. Finally, the loop function $K = K(M_1, M_2)$ is given by

$$K(M_1, M_2) \equiv -\frac{M_1^2 + 4M_1 M_2 + M_2^2}{M_1 M_2 (M_1 + M_2)} + \frac{4(M_1^2 + M_1 M_2 + M_2^2)}{(M_1 - M_2)(M_1 + M_2)^2} \log \frac{M_1}{M_2}. \quad (44)$$

Concerning the θ decay into charged leptons, one needs to add the contribution of the W -exchange diagram. The main phenomenological constraint comes from $\theta \rightarrow e^+ e^-$. Therefore,

in our simplified calculation we identify the active neutrino with the electron neutrino ν_e . Then, the additional contribution to the effective coupling of θ to electrons is given by

$$\mathcal{L}_{\theta e\bar{e}}^W = i \left(\frac{M_{11}}{f} F_1 - \frac{M_{22}}{f} F_2 \right) \frac{2\sqrt{2} G_F}{(4\pi)^2} m_e \bar{e} \gamma_5 e \theta, \quad (45)$$

that carries a relative factor -2 with respect to the Z -exchange contribution. This is why the sum of the two contributions in Eqs. (41) and (45) gives $g_{\theta e\bar{e}}/m_e = g_{\theta u\bar{u}}/m_u = -g_{\theta d\bar{d}}/m_d$, consistently with the one-family result in Eq. (39). Also, the functions $F_{1,2}$ are of the order $m_i^2/M_j \sim m_\nu$, so the couplings are of the same order as in Eq. (39). Still, cancellations between the various terms are possible leading to a suppression of the θ decay width for special flavour structures. This possibility may deserve a future investigation, since an appropriate family symmetry could in principle raise the pNGB lifetime, so that θ could become a viable DM candidate even for $m_\theta > 1$ GeV, a region where freeze-out could lead to the observed relic density. In this paper we do not invoke a family symmetry for the suppression of $g_{\theta f\bar{f}}$ to happen, but rather we rely on the one-family estimate given in Eq. (39).

Gauge bosons. The last effective coupling of the pNGB that may have important phenomenological consequences is the one to photons. Let us discuss first the limit of exact global symmetry, and then comment on the effect of explicit breaking. In general, the couplings of a NGB to gauge bosons are controlled by the gauge anomalies of the associated global symmetry. In the case θ is the Majoron, the global symmetry can be identified with $B-L$, because θ can be rotated away from the Yukawa interactions by rephasing all the SM fermions with a $B-L$ transformation. Since $B-L$ is anomaly-free with respect to the SM gauge symmetries, this rephasing does not generate anomalous couplings of the type $\theta F\tilde{F}$, rather the only leftover interaction is derivative, $(\theta/f)\partial_\mu J_{B-L}^\mu$, which is generated by the redefinition of the fermion kinetic terms. One concludes that the Majoron has no anomalous couplings to gauge bosons.⁸

This discussion is easily generalized to the $U(1)_X$ symmetry that we consider in this paper. First, note that the sterile neutrino charges are irrelevant for gauge anomalies. In order to allow for fermion masses, we can take $X(l_\alpha) = -X(e_\alpha^c) = x_\alpha$ for $\alpha = e, \mu, \tau$, and zero X -charge for the quarks. Then it is easy to check that $U(1)_X$ is anomaly-free with respect to electromagnetism and colour, and thus θ does not couple to photons and gluons. The anomaly with EW interactions is proportional to $\sum_\alpha x_\alpha$ and can also be taken equal to zero for simplicity. In summary, in the limit of exact symmetry our NGB does not couple to gauge bosons.

It is far more difficult to compute the pNGB couplings to gauge bosons in the presence of explicit symmetry breaking sources. One expects that such couplings will arise at some level, since there is no symmetry arguments to prevent them at all orders. However, the

⁸ Equivalently, one may remove θ from the Yukawa interactions by an L transformation; in this case one is left with the couplings $(\theta/f)[\partial_\mu J_L^\mu + (c_W F_W \tilde{F}_W + c_Y F_Y \tilde{F}_Y)/(16\pi^2)]$, where $c_{W,Y} \neq 0$ account for the EW anomalies of the lepton number symmetry. There is no physical difference between the two pictures. In particular one can check that (i) θ does not couple to weak gauge bosons at one-loop even in the L case, because the triangle diagrams generated by the derivative coupling to J_L^μ cancel $c_{W,Y}$ exactly; (ii) θ does not decay to two on-shell quarks even in the $B-L$ case, because the quarks are vector-like under $B-L$.

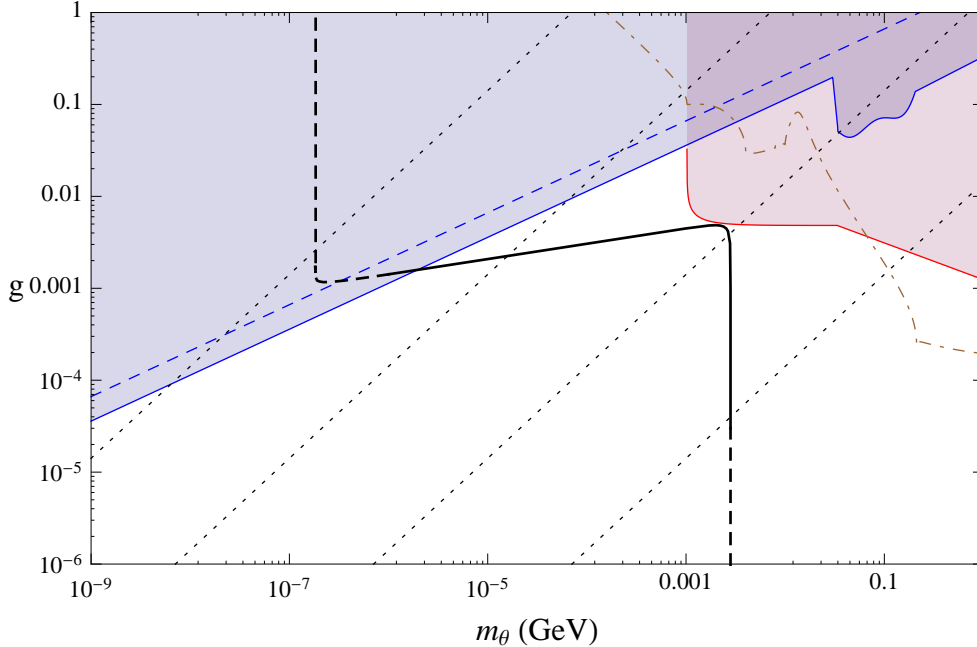


Figure 3: The constraints on the DM lifetime in the $(m_\theta - g)$ plane, for $m_\nu = 0.05$ eV. The thick black curve, corresponding to the correct DM relic density, as well as the dotted lines, corresponding to constant values of m_N , were already presented in Fig. 1. The curve is dashed for $m_\theta \lesssim 1$ keV, because DM is warm in this region (see the text), and for $g \lesssim 10^{-5}$, because of the theoretical bound $f < M_P$, shown in Fig. 1. The blue solid (dashed) line is the upper bound on g from DM decays into neutrinos coming from astrophysics and cosmology (from the universe lifetime): the blue shaded region is therefore excluded. The red solid line is the analog bound for DM decays into e^+e^- : the red shaded region is correspondingly excluded. Finally, the brown dot-dashed line is a conservative estimate of the upper bound on g from DM decays into photons (see the text).

computation of the lowest order non-vanishing contribution is non-trivial. We have shown explicitly above that, in the present framework, θ couples at tree-level to neutrinos only, while, at one-loop, couplings to the Higgs boson as well as to charged fermions are induced. At two-loop there is a number of diagrams that connect θ to two gauge bosons. These diagrams do not necessarily add up to zero: on the one hand, they involve the neutrino mass parameters that break explicitly $U(1)_X$, thus they are not expected to respect the anomaly argument above; on the other hand, however, these parameters may not be sufficient to induce an operator $\theta F\tilde{F}$ already at two-loop order.

The computation of these two-loop diagrams is beyond the purpose of the present paper, because we will show that a decay $\theta \rightarrow \gamma\gamma$ induced at this order would be irrelevant for phenomenology anyway: once the constraints from $\theta \rightarrow f\bar{f}$ are imposed, the surviving DM parameter space is not further reduced by constraints on photons. To see this, we estimate the size of a two-loop contribution by taking the effective one-loop couplings in Eq. (39), and computing the usual fermion triangle diagrams with two final state photons. The effective

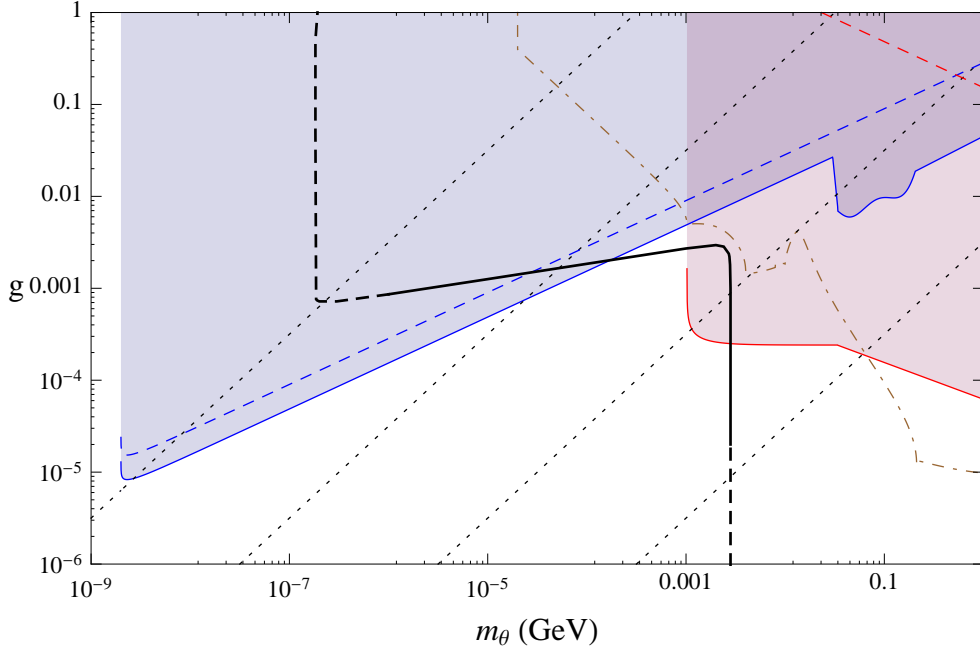


Figure 4: The same as in Fig. 3, but for a heavier neutrino mass scale, $m_\nu = 1$ eV. The extra red dashed line is the upper bound from $1/\Gamma(\theta \rightarrow e^+e^-) > \tau_0$ (in the case of Fig. 3 the analog line lies entirely in the region $g > 1$ and $m_\theta > 1$ GeV, which is not displayed).

lagrangian can be written as

$$\mathcal{L}_{\theta\gamma\gamma} = \frac{1}{8} g_{\theta\gamma\gamma} \theta \epsilon_{\mu\nu\alpha\beta} F^{\mu\nu} F^{\alpha\beta} \quad (46)$$

and the corresponding decay width is

$$\Gamma(\theta \rightarrow \gamma\gamma) = \frac{1}{64\pi} g_{\theta\gamma\gamma}^2 m_\theta^3, \quad g_{\theta\gamma\gamma}^{estimate} \simeq \frac{\alpha}{\pi} \sum_f \frac{g_{\theta f \bar{f}}}{m_f} G\left(\frac{m_\theta}{2m_f}\right), \quad (47)$$

where $G(x) = |\arcsin^2 x|/x^2$ [18]. Note that even in the limit of exact NGB, with $m_\theta \ll m_f$ and thus $G(x) \approx 1$, the sum does not give zero when the sign in Eq. (39) is taken into account. Therefore, we stress again that this is just a conservative estimate, and the computation of all the two-loop contributions may lead to a further cancellation. Neglecting the mass dependence in G , the requirement $\Gamma(\theta \rightarrow \gamma\gamma) < 1/\tau_0$ gives the order of magnitude constraint $g \lesssim 4 \times 10^3 (\text{MeV}/m_\theta)^{3/2} (\text{eV}/m_\nu)$. This bound is weaker than the one obtained below Eq. (40) from $\Gamma(\theta \rightarrow f\bar{f})$, as long as $m_\theta \lesssim 1$ GeV, which is the region relevant in the present scenario.

5.2 Cosmological and astrophysical bounds on θ couplings

In practice, the DM lifetime has to be larger than the lifetime of the universe, because late DM decays affect several cosmological and astrophysical observations. To derive the

corresponding bounds on our scenario we make the assumption that θ is all the DM in the universe. For the sake of simplicity, here too we take the approximation of one lepton family, barring large cancellations between the flavour parameters. A sub-GeV particle could decay into neutrinos, electrons and photons. We discuss these decay channels in turn.

Let us first consider the decay $\theta \rightarrow \nu\nu$. As it was noted in Ref. [5] such a decay could affect the expansion history of the universe, because it represents energy transfer from a non-relativistic θ to relativistic neutrinos. Using SNIa and CMB data, one obtains the bound [19]

$$\Gamma(\theta \rightarrow \nu\nu) < 4.5 \times 10^{-20} \text{ s}^{-1} \quad (\text{or } \tau > 700 \text{ Gyr}) . \quad (48)$$

Using Eq. (37), the corresponding bound on the coupling of θ to neutrinos is

$$g_{\theta\nu\nu} < 4 \times 10^{-20} \left(\frac{\text{MeV}}{m_\theta} \right)^{1/2} . \quad (49)$$

This upper bound can be improved in the range $30 \text{ MeV} < m_\theta < 200 \text{ MeV}$ [20], using searches for the diffuse neutrino supernova background by Super-Kamiokande. For masses $m_\theta > 200 \text{ MeV}$ the best limit come from atmospheric neutrino observations. Since the observed spectrum coincides with theoretical estimates, one can set an upper bound on $g_{\theta\nu\nu}$ [20, 21]. We use the relation in Eq. (38) to translate these constraints in the $(m_\theta - g)$ plane. The blue region shown in Figs. 3,4 is excluded by the combination of the observational bounds discussed above.

Let us now consider the $\theta \rightarrow e^+e^-$ decay. We adapt to our case the analysis performed in Ref. [21]. For $m_\theta \leq 20 \text{ MeV}$, the dominant physical process which constrains the parameter space is annihilation at rest, contributing to the 511 keV line. The limit is approximately given by

$$\left(\frac{\text{MeV}}{m_\theta} \right) \Gamma(\theta \rightarrow e^-e^+) < 5 \times 10^{-27} \text{ s}^{-1} . \quad (50)$$

Using Eq. (40), this limit leads to

$$g_{\theta ee} < 9 \times 10^{-24} \left(1 - \frac{4m_e^2}{m_\theta^2} \right)^{-1/2} . \quad (51)$$

For $20 \text{ MeV} \leq m_\theta \leq 1 \text{ GeV}$ the dominant process is internal bremsstrahlung, i.e., photons radiated from the final electron or positron. One gets a slightly more stringent bound than the one shown in Eq. (50) [21]. The predicted value of the coupling $g_{\theta ee}$ in our model, as follows from Eq. (39), is

$$g_{\theta ee} = g \frac{G_F}{(4\pi)^2} m_\nu m_e = 3 \times 10^{-23} \left(\frac{g}{10^{-3}} \right) \frac{m_\nu}{\text{eV}} . \quad (52)$$

The corresponding observational bound is plotted in Figs. 3,4 as the red shaded region.

Note that the value $m_\theta \simeq 3 \text{ MeV}$, predicted by the Higgs portal, could lead to an excess of 511 keV γ rays from the galactic center of the Milky Way. As discussed in Ref. [22], a $\sim 1.5 \text{ MeV}$ electron or positron, produced by a $\theta \rightarrow e^+e^-$ decay, is enough non-relativistic to subsequently annihilate at rest in the galactic center, leading to a γ line at 511 keV. In

particular, with such a mass one can easily obtain a γ flux excess of the order of the one observed by the INTEGRAL γ observatory [22, 23]. However, unless the DM galactic profile is much more cuspy than the usually considered profiles, a DM decay gives a flux that is not sufficiently peaked around the galactic center [22, 24] to be able to reproduce the morphology of the signal observed by INTEGRAL.

We discuss now the decay $\theta \rightarrow \gamma\gamma$. For small masses, $m_\theta \lesssim \text{keV}$, the photons energy is absorbed by the baryonic gas in the early universe and the processes of recombination and reionization are affected. One can use the analysis done in Ref. [25] (see also [26]) to bound the coupling $g_{\theta\gamma\gamma}$. For larger masses, $m_\theta \gtrsim \text{keV}$, photons are no longer absorbed and propagate freely. Their contribution to the isotropic diffuse photon background allows also to bound $g_{\theta\gamma\gamma}$ [26, 27].

For the region we are most interested in, stronger bounds can be obtained from the gamma-ray line emission limits from the Milky-Way central region. Indeed, in the range $40 \text{ keV} < m_\theta < 16 \text{ MeV}$ [27], one obtains

$$\Gamma(\theta \rightarrow \gamma\gamma) < 10^{-28} \text{ s}^{-1} \left(\frac{m_\theta}{\text{MeV}} \right) . \quad (53)$$

Using Eq. (47), this corresponds to

$$g_{\theta\gamma\gamma} < 4 \times 10^{-21} \left(\frac{\text{MeV}}{m_\theta} \right) \text{ GeV}^{-1} . \quad (54)$$

Taking the estimate for $g_{\theta\gamma\gamma}$ given in Eq. (47), all these constraints translate in an upper bound on g , which is shown in Figs. 3,4 by the brown dot-dashed curve.

We should mention that there are astrophysical constraints on the couplings $g_{\theta\nu\nu}$ [28], $g_{\theta ee}$ [29], and $g_{\theta\gamma\gamma}$ [29, 30], based on stellar energy loss due to θ emission, provided m_θ is low enough to be produced in stellar interiors. These limits are valid without the need to assume that θ is DM. However, these astrophysical limits are in general weaker than the ones mentioned above.

We also would like to note that the region we consider for pNGB DM masses includes the keV scale, corresponding to warm dark matter (WDM). Since at the epoch of structure formation WDM has free-streaming lengths below the Mpc scale, having WDM at least as a non-negligible DM component can alleviate some of the disagreement between the standard cold DM scenario and a variety of galactic observations at small scales [31]. The ideal observation which could place a limit on WDM is Lyman- α forest, i.e. the Lyman- α absorption produced by intervening neutral hydrogen in the spectra of light emitted by distant quasars. Using Ly- α observations together with other cosmological data sets, different groups have put lower limits on m_{WDM} , assuming WDM has a thermal distribution and that it is the whole of DM. The bounds in the literature [32] differ by factors of a few, ranging from 0.5 keV up to 4 keV. In our case, the DM species is not in thermal equilibrium and the bounds can slightly change compared to the WDM thermal relic case. In addition, if WDM is only a part of the whole of DM the bounds are relaxed. We should mention that there are some potential problems in the obtention of these bounds, and they should be regarded as controversial. On the one hand there could be large systematic errors, and on the other hand the Ly- α analysis has to be performed at scales which are already in the non-linear

regime, where calculations are less reliable. If one disposes of the Ly- α data these bounds on WDM disappear altogether [33].

6 Conclusions

We proposed a new pseudo-scalar gauge-singlet DM candidate θ with mass in the keV - MeV range. Its couplings to the SM particles are feeble, because they are mediated by new physics at a large scale f , which we identify with the seesaw scale. The θ relic density can be produced by the scattering with the heavy particles (the sterile neutrinos), at temperatures of the order of f , or alternatively it can be generated at the EW scale through the Higgs portal, by the tiny θ - H coupling $\lambda \sim 10^{-10}$, which is induced by the seesaw interactions. Today, θ decays into light neutrinos and, if it is heavier than one MeV, into e^+e^- , with rates that can saturate the present upper bounds.

We argued that such a candidate is theoretically well-motivated. The heavy new physics sector is generically associated with several global symmetries. Some of the corresponding pNGBs may remain light, if the explicit symmetry breaking effects are sufficiently small. For concreteness, we demonstrated that a $U(1)_X$ family symmetry of the sterile neutrino sector can be broken collectively by a set of neutrino Yukawa couplings, so that the pNGB mass m_θ is proportional to the EW scale and to the (small) product of the Yukawa couplings. In such scenario m_θ is not quadratically sensitive to the cutoff (at leading order), therefore the presence of a light DM scalar below the EW scale is justified.

In order to calculate the θ relic density, we computed the rate of the θ interactions with a sterile neutrino N and, through the Higgs portal, with the SM particles, and we studied numerically the Boltzmann equation for the θ number density. In this framework there are only two independent parameters, the mass of the pNGB, m_θ , and the coupling of the pNGB to the sterile neutrino, g , a feature which makes our scenario especially predictive. We find that the Higgs portal produces the desired DM relic density, through the freeze-in mechanism, for a unique value of the DM mass, $m_\theta \simeq 3$ MeV. This prediction relies only on the relation between the DM mass and its coupling to the Higgs, $m_\theta^2 = \lambda v^2$, which is well-justified in the case of our pNGB. As long as the reheating temperature is sufficiently high, the sterile neutrino portal can produce the required relic density by freeze-in for smaller values of m_θ , from $\simeq 3$ MeV down to 0.15 keV. These results are summarized in Fig. 1.

The constraints from $\Gamma(\theta \rightarrow \nu\nu)$, however, exclude the region $m_\theta \lesssim 1(100)$ keV, for a light neutrino mass $m_\nu = 0.05(1)$ eV. In addition, the constraints from $\Gamma(\theta \rightarrow e^+e^-)$ put an upper bound on the coupling g in the region $m_\theta \gtrsim 1$ MeV: for $m_\nu = 0.05(1)$ eV and $g \sim 10^{-3}(10^{-4})$, the expected electron-positron flux is close to the present sensitivity. These results are summarized in Figs. 3,4.

In turn, the seesaw scale can be constrained, by requiring our pNGB to be a viable DM candidate in agreement with all the bounds above. In the case of freeze-in through the sterile neutrino portal, we find 10^5 GeV $\lesssim m_N \lesssim 10^{10}$ GeV, and a SSB scale $f \sim 10^3 m_N$. In the case of the Higgs portal, one has instead 10^{10} GeV $\lesssim m_N \lesssim 10^{14}$ GeV, and correspondingly 10^{13} GeV $\lesssim f \lesssim M_P$.

Finally, we note that our analysis of the parameter space was mostly performed in the

one-family approximation, that is, neglecting possible hierarchies among the sterile neutrino mass parameters and the neutrino Dirac Yukawa couplings. Some of our main results are independent from this approximation, such as the allowed range for the DM mass. On the contrary, the bounds on g and m_N are obviously sensitive to the flavour structure. A rough idea of these dependence can be grasped by comparing the case $m_\nu = 0.05$ eV (Fig. 3) with the case $m_\nu = 1$ eV (Fig. 4). A more detailed exploration of the flavour parameter space of this scenario will be desirable, in particular, if one wants to compare with neutrino flavour models.

Acknowledgements

It is a pleasure to thank K.S.Babu, X.Chu, J.R.Espinosa, D.Greynat, K.Jedamzik, S.Lavignac, S.Peris, S.Palomares-Ruiz, A.Pomarol, M.Quiros, J.Redondo, J.Serra, M.H.G.Tytgat and A.Varagnolo for very useful discussions. The work of MF was supported in part by the Spanish Research Project CICYT-FEDER-FPA-2008-01430, the Consolider-Ingenio 2010 Programme CPAN (CS-D2007-00042), and the Marie-Curie Reintegration Grant PERG06-GA-2009-256374 within the European Community FP7. MF also thanks the Service de Physique Théorique of the Université Libre de Bruxelles for hospitality. The work of TH is supported by the FNRS-FRS, the IISN and the Belgian Science Policy (IAP VI-11). TH thanks J.R.Espinosa and the IFAE at Barcelona for hospitality. The work of EM has been partly supported by CICYT-FEDER-FPA2008-01430 and by 2009SGR894.

References

- [1] Y. Chikashige, R. N. Mohapatra and R. D. Peccei, Phys. Lett. B **98** (1981) 265.
- [2] E. K. Akhmedov, Z. G. Berezhiani, G. Senjanovic and Z. j. Tao, Phys. Rev. D **47** (1993) 3245 [arXiv:hep-ph/9208230].
- [3] I. Z. Rothstein, K. S. Babu and D. Seckel, Nucl. Phys. B **403** (1993) 725 [arXiv:hep-ph/9301213].
- [4] V. Berezhinsky and J. W. F. Valle, Phys. Lett. B **318** (1993) 360 [arXiv:hep-ph/9309214].
- [5] M. Lattanzi and J. W. F. Valle, Phys. Rev. Lett. **99**, 121301 (2007) [arXiv:0705.2406 [astro-ph]].
- [6] F. Bazzocchi, M. Lattanzi, S. Riemer-Sorensen and J. W. F. Valle, JCAP **0808** (2008) 013 [arXiv:0805.2372 [astro-ph]]; J. N. Esteves, F. R. Joaquim, A. S. Joshipura, J. C. Romao, M. A. Tortola and J. W. F. Valle, Phys. Rev. D **82** (2010) 073008 [arXiv:1007.0898 [hep-ph]].
- [7] P. H. Gu, E. Ma and U. Sarkar, Phys. Lett. B **690** (2010) 145 [arXiv:1004.1919 [hep-ph]].
- [8] L. J. Hall, K. Jedamzik, J. March-Russell and S. M. West, JHEP **1003**, 080 (2010) [arXiv:0911.1120 [hep-ph]].

- [9] C. T. Hill and G. G. Ross, Phys. Lett. B **203** (1988) 125; C. T. Hill and G. G. Ross, Nucl. Phys. B **311** (1988) 253.
- [10] M. Schmaltz and D. Tucker-Smith, Ann. Rev. Nucl. Part. Sci. **55** (2005) 229 [arXiv:hep-ph/0502182]; M. Perelstein, Prog. Part. Nucl. Phys. **58** (2007) 247 [arXiv:hep-ph/0512128].
- [11] W. L. Guo and Y. L. Wu, JHEP **1010** (2010) 083 [arXiv:1006.2518 [hep-ph]].
- [12] M. Farina, D. Pappadopulo and A. Strumia, Phys. Lett. B **688**, 329 (2010) [arXiv:0912.5038 [hep-ph]].
- [13] C. E. Yaguna, JHEP **1108** (2011) 060 [arXiv:1105.1654 [hep-ph]].
- [14] A. Y. Smirnov and R. Zukanovich Funchal, Phys. Rev. **D74**, 013001 (2006) [hep-ph/0603009].
- [15] P. Sikivie, Lect. Notes Phys. **741** (2008) 19 [arXiv:astro-ph/0610440].
- [16] R. Barbieri, L. J. Hall, S. J. Oliver and A. Strumia, Phys. Lett. B **625** (2005) 189 [arXiv:hep-ph/0505124].
- [17] A. Pilaftsis, Phys. Rev. D **49** (1994) 2398 [arXiv:hep-ph/9308258].
- [18] L. Ametller, L. Bergstrom, A. Bramon and E. Masso, Nucl. Phys. B **228** (1983) 301.
- [19] Y. Gong and X. Chen, Phys. Rev. D **77**, 103511 (2008) [arXiv:0802.2296 [astro-ph]].
- [20] S. Palomares-Ruiz, Phys. Lett. B **665**, 50 (2008) [arXiv:0712.1937 [astro-ph]].
- [21] N. F. Bell, A. J. Galea and K. Petraki, Phys. Rev. D **82**, 023514 (2010) [arXiv:1004.1008 [astro-ph.HE]].
- [22] N. Prantzos *et al.*, arXiv:1009.4620 [astro-ph.HE], and Refs. therein.
- [23] G. Weidenspointner, G. Skinner, P. Jean, J. Knodlseder, P. von Ballmoos, G. Bignami, R. Diehl, A. W. Strong *et al.*, Nature **451** (2008) 159-162.
- [24] Y. Ascasibar, P. Jean, C. Boehm and J. Knoedlseder, Mon. Not. Roy. Astron. Soc. **368** (2006) 1695 [arXiv:astro-ph/0507142].
- [25] L. Zhang, X. Chen, M. Kamionkowski, Z. g. Si and Z. Zheng, Phys. Rev. D **76**, 061301 (2007) [arXiv:0704.2444 [astro-ph]].
- [26] X. L. Chen and M. Kamionkowski, Phys. Rev. D **70**, 043502 (2004) [arXiv:astro-ph/0310473].
- [27] H. Yuksel and M. D. Kistler, Phys. Rev. D **78**, 023502 (2008) [arXiv:0711.2906 [astro-ph]].

- [28] J. A. Grifols, E. Masso and S. Peris, Phys. Lett. B **215** (1988) 593.
- [29] G. G. Raffelt, “Stars As Laboratories For Fundamental Physics: The Astrophysics Of Neutrinos, Axions, And Other Weakly Interacting Particles,” *Chicago, USA: Univ. Pr. (1996) 664 p*
- [30] E. Masso and R. Toldra, Phys. Rev. D **52** (1995) 1755 [arXiv:hep-ph/9503293].
- [31] B. Moore, S. Ghigna, F. Governato, G. Lake, T. R. Quinn, J. Stadel and P. Tozzi, Astrophys. J. **524**, L19 (1999); G. Gentile, A. Burkert, P. Salucci, U. Klein and F. Walter, Astrophys. J. **634**, L145 (2005) [arXiv:astro-ph/0510607].
- [32] M. Viel, J. Lesgourgues, M. G. Haehnelt, S. Matarrese and A. Riotto, Phys. Rev. D **71**, 063534 (2005) [arXiv:astro-ph/0501562]; K. Abazajian, Phys. Rev. D **73**, 063513 (2006) [arXiv:astro-ph/0512631]; M. Viel, J. Lesgourgues, M. G. Haehnelt, S. Matarrese and A. Riotto, Phys. Rev. Lett. **97**, 071301 (2006) [arXiv:astro-ph/0605706]; U. Seljak, A. Makarov, P. McDonald and H. Trac, Phys. Rev. Lett. **97**, 191303 (2006) [arXiv:astro-ph/0602430]; M. Viel, G. D. Becker, J. S. Bolton, M. G. Haehnelt, M. Rauch and W. L. W. Sargent, Phys. Rev. Lett. **100**, 041304 (2008) [arXiv:0709.0131 [astro-ph]].
- [33] S. Hannestad and G. Raffelt, JCAP **0404**, 008 (2004) [arXiv:hep-ph/0312154].